The Choice Architecture of Personal Budgeting Tools: Deconstructing the Effects of Budgeting Partitioning

Colin West[1], Gulden Ulkumen[2], Patrycja Arundel[3], and Craig R. Fox[4]

[1]University of Toronto, Rotman School of Management

[2]University of Southern California, Marshall School of Business

[3]University of California, Irvine, Paul Merage School of Business

[4]University of California, Los Angeles, Anderson School of Management

WORKING PAPER
May 2022

ABSTRACT

Most people want to save more money, but they are often unsure how much they can or should allocate towards savings at any given time. To simplify this decision, people may rely on a household budget, organizing their finances into categories using spreadsheets, online banking software, mobile applications, and other budgeting tools. Across eight preregistered experiments, we show that the configuration of these budgeting tools can have large, unintended effects on saving versus spending decisions. In Studies 1-4, we find evidence that financial allocation decisions can be biased by how budget categories are partitioned, consistent with a 1/n heuristic. In Studies 5-7, we identify the specific features of budgeting tools that lead people to rely on a 1/n heuristic, to a greater or lesser extent. Finally, in Study 8, we combine these insights to design a budgeting tool that can nudge people to allocate more money towards saving. We conclude by discussing the implications for personal finance and other resource allocation decisions.

Abstract word count: 162

Full manuscript word count:

Keywords: financial decision-making, budgeting, saving, heuristics, biases, choice architecture, partition dependence.

Most people recognize that saving money is critical for their financial security, but they are unsure about how much to save at any given time. Deciding how much to save involves a complex trade-off between financial obligations, immediate consumption, and long-term plans as well as predictions about future earnings and expenses. To simplify this decision, people use a wide array of budgeting tools. For instance, people use spreadsheets, online banking software, and mobile fintech applications to divide their household budget into meaningful categories so that they can more easily decide how to allocate their paycheck and any windfalls (e.g., a holiday bonus, sales commission, tax refund, etc.). In lieu of such digital tools, people budget by physically separating cash into multiple envelopes, giving cash to friends or family members for safe keeping, purchasing illiquid assets as a means of protecting and storing wealth, opening multiple bank accounts, joining informal savings clubs, and using handwritten notes, among other methods (Zelizer, 1989; 2017). Indeed, people across cultures and generations have relied on budgeting tools to organize their financial priorities and guide their allocation decisions (Graeber, 2011). All of these disparate budgeting tools are *intended* to help people make better financial decisions, but they can also lead to predictable biases depending on how they are designed.

In this research, we identify some essential design features of budgeting tools that can influence users' financial allocation decisions. Across eight preregistered experiments (N = 4,860), we show that the configuration of one's household budget can act as *choice architecture* – shaping the decision process that people use to allocate money and, ultimately, the amount of money people save versus spend. Furthermore, we identify the specific features the lead some budgeting configurations to encourage more saving than others. Specifically, we demonstrate that when decision-makers use certain budgeting configurations, they are biased in the direction of spreading money evenly over identified financial categories. This can lead to a *partition dependence* bias such that their chosen allocation varies systematically depending on the arbitrary partitioning of their budget. To illustrate this phenomenon, consider the decision of allocating a holiday bonus across two budget categories: 1) spending and 2) savings. Many people apply a *1/n heuristic*, contributing 50% of their bonus to spending and 50% to savings. Now instead, consider a budget in which the savings category is partitioned into discrete sub-categories such that there are a total four budget categories: 1) spending, 2) saving for emergencies, 3) saving for upcoming expenses, and 4) saving for retirement. Applying the same 1/n heuristic would lead people to allocate 75% of their bonus to savings.

In Studies 1-4, we find evidence indicating that people do indeed rely on a 1/n heuristic when using certain types of budgeting tools, but not others. Furthermore, we show that budget partitioning can have a stronger impact on savings decisions than setting savings goals. We observe a robust effect of budgeting partitioning when people are allocating their ordinary income, a raise, or a bonus.

In Studies 5-7, we identify the specific features of budgeting tools that lead people to rely on a 1/n heuristic in their financial allocation decisions. We find that people are more likely to use a 1/n heuristic under specific budget configurations and, therefore, setting more savings goals or creating more earmarked savings accounts will not necessarily lead to an increase in overall savings. Indeed, setting too many savings goals can backfire, leading to a decrease in overall savings because people often feel demotivated when they have multiple competing goals (Soman & Zhao, 2011). Therefore, we sought to explain why setting multiple savings goals can *decrease* overall savings when people use some budgeting tools and *increase* overall savings when people use other tools.

Lastly, in Study 8, we combine these insights to design and test a budgeting tool that can nudge people to increase their total allocation to savings. Here, we find that a treatment budgeting tool can increase the amount of money that people allocate to savings by 23-percentage-points, relative to control condition in which people were prompted to set savings goals. In fact, we found that prompting people to set savings goals had no effect on subsequent savings decisions unless people used a budgeting tool that explicitly partitioned their goals during the allocation process. This is a striking result given that 'set savings goals' is often the primary financial advice offered by banks, wealth managers, and fintech software platforms. In contrast, the financial advice industry has paid much less attention to the configuration of the financial tools that people use to guide their decision-making.

Taken together, the experiments presented here show that budget configuration can act as a powerful tool of choice architecture and play a critical role in shaping household financial allocation decisions.

Background and hypotheses

The role of budgeting tools in household financial decision-making

The primary purpose of budgeting tools is to help people resist temptations and stick to their long-term financial plans. Many people are well-aware that they are susceptible to short-sighted financial decision-making (Laibson, 1997; O'Donoghue & Rabin, 1999), impulsive spending (Hoch & Loewenstein, 1991; Fugita, 2011), and social pressures to overspend on conspicuous goods (Bagwell & Bernheim, 1996; Charles et al., 2009). We turn to budgeting tools to resist these tendencies and make financial decisions that will improve our long-term financial wellbeing.

While there are many different types of budgeting tools – from simple handwritten ledgers to mobile fintech applications – there are a few essential features that they all share. All budgeting tools involve partitioning one's finances into categories and then following a set of explicit or implicit rules to decide how much money to allocate to and spend from each category. Where budgeting tools vary is in the granularity and partitioning of financial categories as well as in the procedures for allocating and spending money (Zhang et al., 2020). For instance, some budgeting tools include just two categories: spending and savings. Others might partition these over-arching categories into multiple sub-categories for spending (e.g., groceries, dining out, clothes shopping, internet, transportation, etc.) and saving (e.g., emergency saving, retirement savings, saving for a new car, saving for a vacation, etc.). Furthermore, some tools prompt people to allocate money across their budget *simultaneously*, for instance, by earmarking money at the beginning of each month or right after receiving a windfall. Other budgeting tools encourage a more *sequential* approach, whereby people first pay their critical bills, then decide how much to spend on discretionary purchases, and then allocate any leftover money towards savings. In this research, we argue that such design features are critical because they can shape the decision-making process that people use to allocate money.

Past research on budgeting tools has focused on their role in curbing impulsive spending. A large body of evidence has shown that categorizing and earmarking money are indeed effective strategies to reduce temptation spending (Shefrin & Thaler, 1981; 2004; Henderson & Peterson, 1992; Wertenbroch, 2002; Benabou & Tirole, 2004; Antonides et al., 2011; Beshears et al., 2016; Galperti, 2019; Koszegi & Matejka, 2020). For instance, consumers can strategically partition their budget, knowing that they will be less likely to spend money that has been explicitly earmarked for saving. Soman and Cheema (2011) tested this budgeting strategy in an experiment with construction workers in rural India. The workers opted-in to receive advice from financial planners who encouraged them to save 40 rupees per week out of their weekly salary of 670 rupees (a 6% savings rate). In the control condition, participants received their full weekly income (670 rupees) as cash in a single envelope. In the treatment condition, participants received the same weekly salary, but split into two separate, sealed envelopes: one with 40 rupees earmarked for saving and another with the remaining 630 rupees. The results showed that people were less likely to spend from the partitioned 'savings' envelopes, leading to greater overall savings accumulation in the treatment group. This experiment demonstrated that partitioning income can be an effective strategy in reducing spending.

However, there is limited research examining a distinct function of budgeting tools – their role in initial allocation decisions. To the best of our knowledge, there is no empirical research examining the effects of budget configuration on allocation decisions (e.g., in the study by Soman and Cheema, the weekly amount allocated towards savings each week was determined by the experimenters). Initial financial allocation decisions are critical since people tend to treat money as non-fungible once it has been earmarked for a specific purpose (Heath & Soll; 1996; Hastings & Shapiro, 2013; 2018; Lian, 2020).

How budgeting tools shape financial allocation decisions

Whenever people receive a sum of money, be it their monthly paycheck or an unexpected windfall, they must decide how to allocate it across their financial priorities. Rather than weighing all the complex trade-offs and consequences involved in this decision, people often rely on shortcuts such as defaults, reference points, and heuristics. For example, many people simply stick with the default savings contribution rate defined by their workplace 401(k) program (3-6% of one's salary are common default savings rates; Thaler & Benartzi, 2004). Similarly, people might take cues from their peer group, using the spending and savings habits of their co-workers, friends, or family members as a reference point for their own allocation decisions (Lindbeck, 1997; Bursztyn et al., 2014). Another shortcut that people might use, under specific conditions, is the naïve diversification heuristic, also referred to as the 1/n heuristic (Messick, 1993). Applying the 1/n heuristic entails the following decision process: start with equal division of resources across identified categories, and then adjust based on the details of the situation (Messick & Schell, 1992). This heuristic is useful and efficient in many types of allocation decisions – it effectively reduces cognitive complexity, and it is easy to explain and justify to others (Messick 1993; Samuelson & Allison, 1994; Fehr & Schmidt, 1999; Kőszegi & Matějka, 2020). For instance, returning to our example of deciding how to allocate a holiday bonus, the 1/n heuristic can be an efficient way to simplify this decision. A decision-maker might want to divide their bonus across discretionary spending, gifts, paying down debts, investment contributions, and emergency savings. If they considered all the relevant information, determining the optimal allocation would be extremely complex and effortful. For example, relevant information could include the expected return on investment contributions, the interest rate on their outstanding debts, the costs of various goods they might want to purchase and predictions about their future earnings and expenses (i.e., to help determine the amount of emergency savings they need). Using a 1/n heuristic allows decision-makers to reduce the amount of information they consider and simplify how they use this information. Equal division (1/n) is treated as a benchmark and relevant information is only considered when making adjustments from this benchmark. While a 1/n heuristic is useful and efficient, it can also lead to systematic biases because people tend to make insufficient adjustments from the benchmark of equal division, leading to resource allocation decisions that are *partition-dependent*, such that people are biased by the partitioning of the option space (Fox & Rottenstreich, 2003; Fox et al., 2005; Langer & Fox, 2005; Tannenbaum et al., 2014; Bardolet et al., 2011).

Previous research has found that people will sometimes rely on a 1/n heuristic in resources allocation decisions including purchasing choices (Roch et al., 2000; Fox, Bardolet, & Leib, 2005), investments (Benartzi & Thaler, 2001; Hedesstrom, Svedsater, & Garling, 2009), and charitable giving (Jung, Nelson, Gneezy, & Gneezy, 2017). 1/n heuristics have also been documented in organizational decision-making including companies' distribution of capital across departments and profits among employees (Harris & Joyce, 1980; Scharfstein & Stein, 2002; Bardolet, Fox, & Lavallo, 2011; Xing et al., 2020). However, the current paper is the first to document this phenomenon in the domain of household financial decision-making. More importantly, we identify the specific conditions and features of budgeting tools that lead people to be more likely to rely on a 1/n heuristic as opposed to the many other simplifying strategies that people use when making financial allocation decisions (e.g., defaults, reference points, or other heuristics).

Hypotheses, predictions, and overview of studies

In this research, we examine how different budgeting tools can shape financial allocation decisions. We theorize that people will be more likely to rely on a 1/n heuristic when using certain budgeting procedures and, therefore, people can be nudged to increase their total saving allocation via a partitioning intervention. Specifically, we predict that people will be more likely to rely on a 1/n heuristic given three related features of the budgeting tool they use.

First, people will be more likely to rely on a 1/n heuristic when they use a budgeting procedure in which they allocate an explicit and fixed sum of money across their budget (e.g., allocating their entire $4,000 monthly income, or a $1,000 bonus), as opposed to deciding how much to save without this fixed reference amount (e.g., “how much money do you want to allocate to savings right now”).

Second, people will be more likely to rely on a 1/n heuristic when they use a budgeting tool that includes an exhaustive set of categories, such that people allocate the entire fixed sum of money with no default category or implied outside categories. Most budgeting tools (including most online banking platforms) prompt people to earmark only a portion of incoming funds for specific purposes and to leave the remainder in their checking account (i.e., making their checking account the de facto default category). We predict that people will be more likely to use on a 1/n heuristic when they are prompted to earmark every dollar with no explicit or implied default category for money not earmarked for a specific purpose.

Third, people will be more likely to rely on a 1/n heuristic when they use a budgeting tool that encourages them to allocate money across all budget categories simultaneously, versus sequentially. For instance, when they divide a sum of money between saving and spending at the same time (e.g., on a single screen), rather than considering each budget category in isolation, paying particular bills and expenses ahead of making savings decisions, or vice versa.

Each of these features of budgeting choice architecture encourage people to consider their whole budget, rather than isolated trade-offs which is how people often approach resource allocation decisions. That is, people tend to focus on the costs and benefits of a single purchase, bill, or savings contribution without considering the full set of options or opportunity costs (Read & Loewenstein, 1995; Fredrick et al., 2009; Hsee & Zhang, 2010) Budgeting tools that encourage people to consider their whole budget may lead to more optimal decision-making (Hsee et al., 1999; Mogilner, Shiv, & Iyengar 2013; Basu & Savani, 2017; 2019), but as we demonstrate, these features can also make people more reliant on a 1/n heuristic and more susceptible to biased decision-making depending on the partitioning of budget categories. Therefore, it is critical that we develop a better understanding of partition-dependence in household budgeting, including understanding how specific features of budgeting tools can influence the decision process that people use to make financial allocation decisions.

We tested our predictions across eight pre-registered studies. First, we tested the existence of a partition dependence bias in household budgeting. In Study 1, we examined how different budget configurations influenced intentions to save versus spend. Participants allocated a hypothetical raise across their household budget with either spending or saving categories partitioned into multiple sub-categories. In Study 2, we conducted a conceptual replication to test the effects of budgeting partitioning on a savings decision with real stakes. Participants allocated a $1000 cash prize towards spending and saving, and we manipulated whether the savings category was partitioned into one overall category versus four sub-categories. We found evidence of a large partition dependence bias, demonstrating that consumers' saving decisions can be highly malleable depending on the arbitrary partitioning of financial categories in a budgeting tool.

Second, we sought to distinguish the effects of budget partitioning from financial goal setting. Typically, when savings are partitioned into multiple sub-categories (e.g., in online banking software, fintech applications, and spreadsheet-based budgeting tools), these categories or accounts are labelled according to savings goals. That is, the effects of partitioning and goal setting are usually confounded. Similarly, in studies 1 and 2 when we partitioned savings into multiple sub-categories, the observed increase in total savings could have been caused either by a 1/n heuristic or by the act of setting savings goals, drawing attention to multiple reasons for savings, or otherwise signaling the importance of savings. Therefore, in Study 3, we isolated the effects of goal setting from partitioning effects. We found that setting saving goals had no effect on participants' subsequent financial allocation decision unless their chosen savings goals were partitioned into distinct categories. In Study 4, we further explored the decision process by varying the number of partitioned savings categories (one, four, or seven categories). Consistent with a tendency to rely on a 1/n heuristic, partitioning savings into more sub-categories led to greater overall savings.

Third, we examined the specific features of budgeting tools that make people more likely to rely on 1/n heuristic. In Studies 5-7, we tested our predictions about the three features of budgeting tools that would lead to a greater reliance on a 1/n heuristic, and therefore make people more susceptible to a partition dependence bias: 1) explicit and fixed sum allocation, 2) exhaustive budget categories, and 3) simultaneous allocation.

Finally, in Study 8, we combined the insights from the preceding studies to measure the relative effects of each budgeting feature on allocations to saving. Relative to a control budget, we found that a treatment budget combining these insights led to a 41% (d = 1.06) increase in total allocation savings.

H1: financial allocation decisions can be influenced by the arbitrary partitioning of categories in a household budgeting tool

H2: partitioning effects are driven by the configuration of the budgeting tool and not by directing attention to savings goals.

H3: partitioning effects are stronger (i.e., people are more likely to rely on a 1/n heuristic) when people use a budgeting tool in which they allocate an explicit, fixed sum

H4: partitioning effects are stronger when people use a budgeting tool that includes an exhaustive set of budget categories.

H5: partitioning effects are stronger when people use a budgeting tool that promotes *simultaneous* (versus sequential) allocation across all identified budget categories.

Table 1. Overview of studies

Study	**Hypothesis tested**	**Money allocated**	**Experimental manipulation**
Study 1 (*N* = 412): The effects of partitioning savings versus spending on financial allocation decisions	H1	Hypothetical 20% salary raise	Participants allocate a hypothetical **salary raise** into their household budget. We manipulated the partitioning of the savings and spending categories.
Study 2 (*N* = 930): Conceptual replication with real stakes	H1	$1,000 cash prize (with an incentive compatible lottery)	A limitation of the previous study was that participants allocated a hypothetical raise. In this study, participants allocated a **real $1,000 cash prize** towards spending (money sent immediately) and saving (money sent in 6 months, plus 10% interest). We manipulated the partitioning of savings into 1 versus 4 categories.
Study 3 (*N* = 239): The effects of budget partitioning versus setting savings goals	H2	Hypothetical 20% holiday bonus	Participants allocated a hypothetical **bonus** into their household budget. We separated the effects of budget partitioning versus goal-setting on allocations decisions. Participants select 3 savings goals (from a list of 12), and we manipulated whether their budget was partitioned according to their selected goals
Study 4 (*N* = 316): The effects of additional savings categories	H2	Hypothetical 20% salary raise	Participants allocated a hypothetical **raise** into their household budget. We manipulated whether savings was partitioned into 2, 4, or 11 sub-categories. If people do, indeed, rely on a 1/n heuristic, then more savings sub-categories should lead to greater total savings allocations.
Study 5 (*N* = 948): Fixed sum budget allocation	H3	Typical monthly income	Participants set up a hypothetical **auto-deposit** towards their savings. They decided how much money they would like to automatically contribute to saving each month. We manipulated the partitioning of savings categories and the salience of their monthly income (fixed sum) at the time of the allocation decision.
Study 6 (*N* = 681): Exhaustive (vs. non-exhaustive) budget categories	H4	Hypothetical 20% salary raise	Participants allocate a hypothetical **raise** into their household budget. We manipulated the partitioning of savings categories and whether a spending category was included, such that the identified budget categories were exhaustive.
Study 7 (*N* = 312): The effects of budgeting simultaneously (vs. sequentially)	H5	Typical monthly income	Participants constructed a realistic **household budget** for the coming months, allocating their expected average monthly income. We manipulated whether they did so sequentially (allocating to each budget category one at a time, on separate pages) versus sequentially (all budget categories on a single page).
Study 8 (*N* = 1022): Combining each feature of budgeting	Combined	Hypothetical 20% holiday bonus	Participants allocated a hypothetical **bonus** into their household budget. We combined the insights from all of the previous studies, measuring the relative effects of each budgeting feature on allocations to saving.

STUDY 1: the effects of budgeting partitioning on savings decisions

In the first study, we tested whether people display partition-dependent preferences in household budget allocation decisions. We randomly assigned participants to make an identical financial decision – allocating a hypothetical raise – using one of three budgeting tools. We measured whether the partitioning of their assigned budgeting tool influenced the amount of money they chose to allocate towards saving.

Method

Study 1 was a preregistered experiment conducted with a sample of 412 online participants (M_{age} = 34.3, *SD* = 10.0, 38% women; M_{income} = $47,891, *SD* = $27,512; 77% employed full-time). In order to be eligible to participate, individuals had to report an annual income greater than USD $10,000 and less than USD $500,000. After providing their personal annual income (in dollars, before taxes), participants received the following instructions: "Imagine you received a 20% raise on your annual income. This amounts to a raise of $[reported annual income * 0.2]. Being as realistic as possible, please indicate how you would allocate this raise into your household budget." Participants were randomly assigned between-subjects to one of three elicitation procedures to decide how much money (in dollars) they would allocate to each budget category. In the 'spending-partitioned' condition, participants were asked to allocate their raise among 7 spending categories (food and dining; housing, repairs, purchases for the household; shopping and personal care; transportation, travel; health, fitness; entertainment products, events; all other spending) and 1 superordinate savings category. In the 'savings-partitioned' condition, participants allocated money across 4 savings categories (i.e., emergency savings, savings for upcoming expenses or purchases, retirement savings, and all other savings) and 1 superordinate spending category. In the control condition, participants allocated money across between 1 superordinate spending category and 1 superordinate savings category.

We held information constant across conditions by listing all of the spending and savings sub-categories in parentheses alongside the superordinate category. For example, in the spending-partitioned condition, the superordinate savings category was presented as follows: *"Savings (emergency savings, savings for upcoming expenses or purchases, retirement savings, and all other savings)."* Additionally, we counter-balanced the order of saving and spending categories in all conditions. See Supplemental Materials for more detail on experimental stimuli.

As the key dependent variable, we measured the amount of money participants allocated towards savings as a percentage of their hypothetical raise. In the spending-partitioned and control conditions, this was defined as a percentage of the raise allocated to the superordinate savings category. In the savings-partitioned condition, this was defined as the sum allocated to emergency savings, savings for upcoming expenses or purchases, retirement savings, and all other savings as a percentage of the hypothetical raise.

Results and discussion

Participants allocated 60.49% (SD = 24.01%) of their raise to savings in the control condition, 41.61% (SD = 28.04%) in the 'spending-partitioned' condition, and 78.83% (SD = 18.02%) in the 'savings-partitioned' condition, $F(2, 409) = 87.75$, $p < .001$. In planned pairwise comparisons, we observed a significant increase in the savings-partitioned condition relative to the control condition, $F(1, 409) = 40.9$, $p < .001$, and relative to the spending-partitioned condition, $F(1, 409) = 165.5$, $p < .001$. We also observed a significant decrease in the spending-partitioned condition relative to the control, $F(1, 409) = 43.9$, $p < .001$. See Figure 1.

As indicated in the preregistered analysis plan, we conducted a follow-up ANCOVA to test whether the results hold controlling for annual income (log), age, gender, and education level. After including these controls, we observed no substantive change in the pattern of results (omnibus ANCOVA: $F(6, 403)=80.8$, $p<.001$). We also found no interaction effect between condition and income (log), suggesting the people across the income spectrum can exhibit partition-dependent preferences in their financial allocation decisions. See Table SMX for robustness checks.

Figure 1. Study 1: effects of condition on percentage of raise allocated to savings.

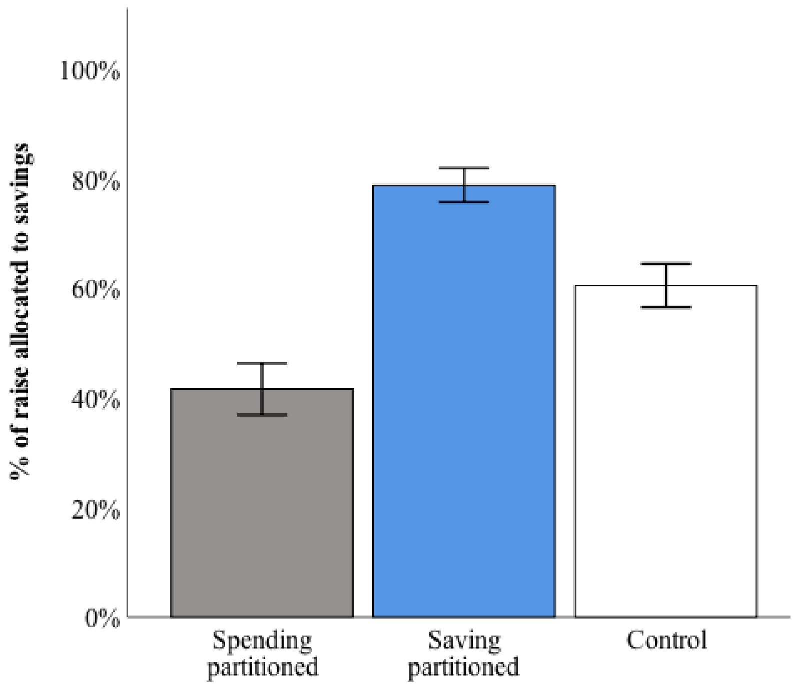

Notes. Displaying means and 95% confidence intervals.

This study provides evidence that intentions to save (versus spend) can be influenced by the arbitrary partitioning of household budget categories. We found large differences in the amount of money people allocated towards saving across our three conditions, which approximate three realistic budgeting tools that consumers might use to help them make a financial allocation decision. The pattern of results is consistent with participants relying on a 1/n heuristic.

STUDY 2: a conceptual replication with real stakes

In Study 2, we conducted a conceptual replication with an incentive compatible measure of savings versus spending. One limitation of the previous study was that participants allocated a *hypothetical* raise. It is possible that people will make different financial choices when there are real stakes. That is, people may be less susceptible to a partition dependence bias when they are making allocation decisions with real financial consequences. To examine this possibility, we collected a representative sample of US adults and offered them the chance to receive a large sum of money immediately or set aside a portion of this money towards 'savings' (to be received in 6 months, plus interest). Participants were randomly assigned to one of two budgeting tools to make this decision.

Method

Study 2 was a preregistered experiment conducted with 930 online participants in the United States (M_{age} = 43.82, *SD* = 15.85, 56% women; median annual income = $50,000; 60% employed full-time). We used a sampled participants stratifying on income to ensure that we recruited participants from across the income distribution, including over-sampling individuals who earned less than $40,000 such that this group comprised at least one third of the overall sample.

Participants enrolled in this study were offered the chance to receive a real $1,000 cash prize via a check in the mail. They could choose to receive the full sum immediately (to be sent in a check within 2 days) or set aside a portion of this money towards 'savings.' Participants were informed that any money allocated towards saving would be sent by mail in a separate check in 6 months, plus 10% interest. We explained that one person from this study would be randomly selected to receive this money for real: *"If you are selected, you will be asked to provide your mailing address so that we can send you two checks in the mail. The 1st check will be for the amount you allocate to spending, and it will be mailed within 2 days. The 2nd check will be for the amount you allocate to savings (plus 10% interest), and it will be mailed in 6 months."* We also provided three examples to ensure that participants understood the decision they were about to make: *"If you allocate all of the money to spending, you will receive $1000 in 2 days and nothing in 6 months; If you allocate half of the money to spending and half to savings, you will receive $500 in 2 days and also $550 in 6 months; If you allocate all of the money to savings, you will receive nothing in 2 days and $1100 in 6 months)."* Therefore, this windfall allocation task was incentive compatible such that participants were aware that their decision had real economic consequences.[1] The interest rate of 10% was chosen based on the results of a pilot experiment which showed that this rate was sufficient to incentivize most participants to allocate a portion of the cash prize to savings.

In order to make this decision, participants were randomly assigned between-subjects to one of two elicitation procedures. In the control condition, participants allocated the $1,000 prize across one spending category (amount they wish to put on the immediate check) and one savings category (amount they wish to put on the check in 6 months, plus 10% interest). In the treatment condition, participants allocated the prize across one spending category and 6 savings categories (with summed amount to be put on the check in 6 months, plus 10% interest). We held

[1] Past research has demonstrated that paying one randomly selected participant in a decision task is an effective incentive compatible mechanism that produces statistically indistinguishable results from equivalent tasks in which all participants receive incentives with no element of chance (Cubitt et al., 1998; Azrieli et al., 2018; 2020). Therefore, we would expect to observe a similar allocation pattern if we were to provide all participants with a $1,000 windfall, rather than using a lottery to select one participant.

information constant across conditions by listing common spending and savings categories in the control condition, which were then partitioned in the treatment condition (see Figure 2).

Figure 2. Study 2: control and treatment budget procedures to allocate a $1,000 cash prize

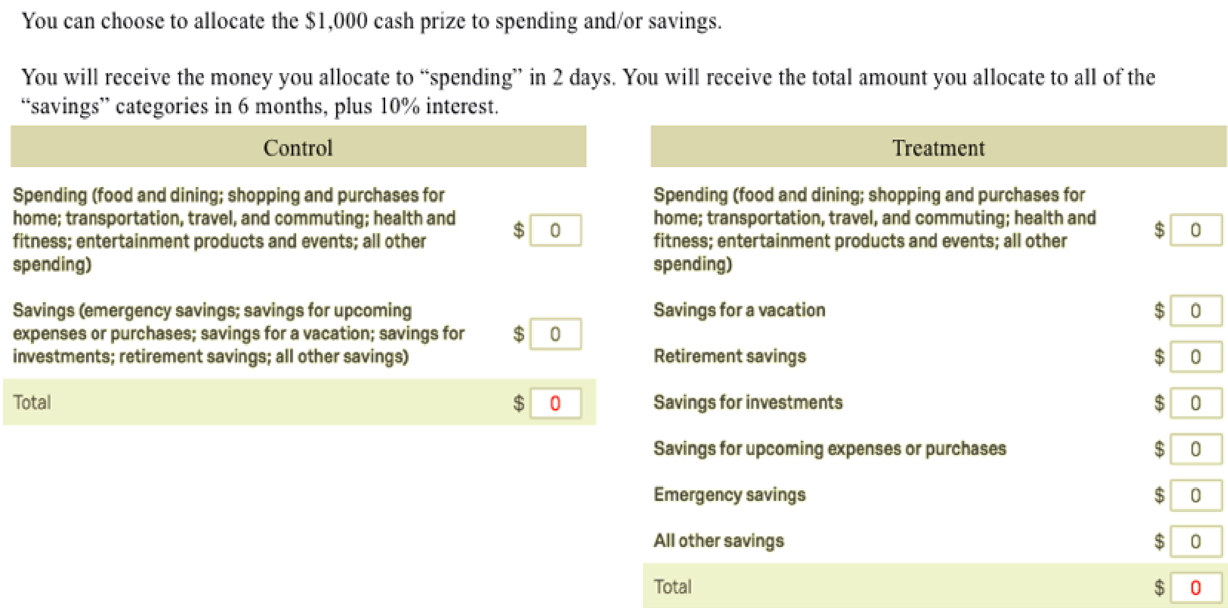

Notes. Example stimuli from Study 2. Participants' responses had to sum to $1,000 in each condition. Order of spending and savings categories was counterbalanced.

Results and discussion

The results of a preregistered ANCOVA showed that, controlling for annual income(log), participants in the treatment condition allocated an extra $179 to savings, compared to those in the control condition ($M_{\text{treatment}}$ = $716.76, SD = 285.99; M_{control} = $558.13, SD = 288.13; F(1, 929) = 70.96, p<.001). See Figure 3. As predicted, these results hold in a regression model controlling for age, education level, and financial literacy (See Table X). Also, we observed no interaction effect between condition and annual income(log), b = 17.94, t(926) = .91 p = .363, CI(b) = [-20.72, 56.59].

In a follow-up exploratory analysis, we examine the role of participants' financial literacy. We measured financial literacy using five questions developed for the Federal Reserve

Survey on Household Economics and Decision-Making (2019), for example: "suppose you had $100 in a savings account and the interest rate was 2% per year. After 5 years, how much money would you have in the account if you left the money to grow?" (More than $102, Exactly $102, Less than $102). We observed a significant main effect of financial literacy such that participants who scored higher on financial literacy contributed more to savings, b = 25.53, $t(928) = 3.72$ $p <$ .001, CI(b) = [12.122, 39.15]. We also observed a significant interaction effect between condition and financial literacy on amount saved, b = -42.90, $t(926) = -3.24$ p = .001, CI(b) = [-68.88, -16.93]. A bootstrapped moderation analysis shows that people at all levels of financial literacy exhibited partition-dependence, but partitioning effects were stronger for those who were low financial literacy (see SM for results of this moderation analysis).

Figure 3. Study 2: effects of condition on the portion of the $1,000 cash prize allocated to savings (check in 6 months, plus 10% interest) versus spending (check in 2 days).

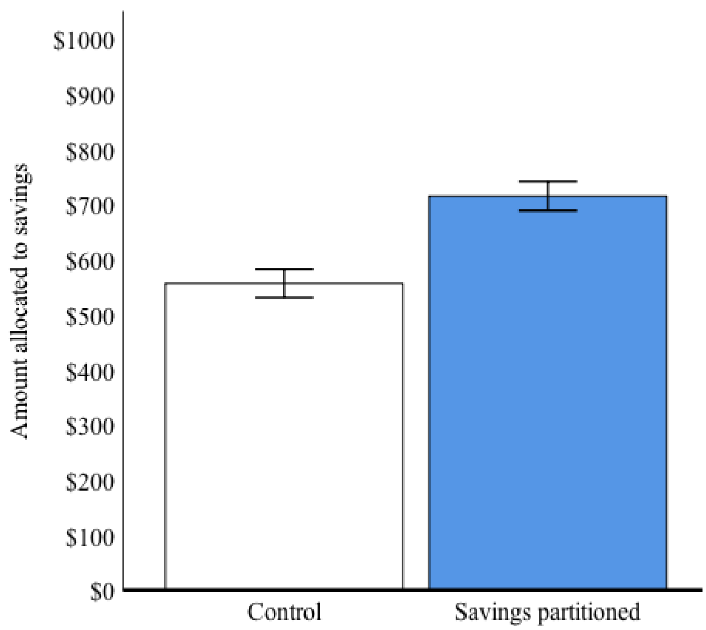

Notes. Reporting marginal means, controlling for annual income (log) and 95% confidence intervals.

The results of this study directly contradict the axiom of *procedure invariance* in rational economic theory. The treatment and control conditions involved an identical decision between a smaller-sooner reward and the option to receive a larger amount in 6 months. All information was held constant, the only difference between conditions was the procedure through which we elicited participants' preference. Taken together, Studies 1 and 2 provide initial evidence that people will rely on a 1/n heuristic under specific budgeting conditions and, therefore, partitioning budget categories can have a large influence on both intentions to save and savings decisions with real stakes.

STUDY 3: the effects of savings goals versus budget partitioning

In this study, we examined the effects of setting savings goals versus budget partitioning on financial allocation decisions. In the two previous studies, we proposed that the observed effects of budget partitioning on savings decisions were driven by the extent to which participants relied on a 1/n heuristic to guide their decision process. However, a plausible alternative explanation is that partitioning savings into sub-categories directed participants' attention to multiple reasons for saving or reminded them of reasons for saving that they may not have otherwise considered. If this alternative explanation is true, then people will allocate more money to savings when they are prompted to set savings goals or when their savings goals are made more salient. In Study 3, we directly tested this attention-based explanation against our proposed heuristic-based explanation.

In addition to providing insights on the mechanism, Study 3 also allows us to measure the relative effects of goal-setting versus budgeting procedure in financial allocation decisions. Goal-setting and budgeting are two fundamental and related strategies that people use to guide their financial decisions (Shefrin & Thaler, 1988; Soman & Cheema, 2004; Fry et al., 2008; Greenberg & Hershfield, 2018; Zhang et al., 2020). However, it is difficult to distinguish their independent effects since financial planning tools typically involve both setting savings goals and then creating budgets in pursuit of those goals. To the best of our knowledge, no existing research has isolated the distinct effects of goal-setting from budget configuration.

Method

Study 3 was a preregistered experiment conducted with 239 online participants in the United States (M_{age} =37.27, *SD* = 11.87; 37% women; median annual income = $36,000; 79% employed full-time). After they provided their monthly income, we asked participants to allocate a hypothetical 20% holiday bonus into their household budget. We calculated each participants' bonus amount (in dollars) and presented them with one of three budgeting tools to decide on how to allocate this money. In the control condition, participants indicated how much of this bonus they would allocate to savings (using a single savings account) versus spending (using a single checking account). In the 'goals' condition, participants first selected three savings goals from a list of 12 common goals (e.g., safety net, retirement, education, new home, etc.), then chose how much of their bonus to allocate into one overall savings account (with their chosen savings goals listed in parentheses) and one spending account. In the 'goals+partitioning' condition, participants selected three savings goals, then chose how much of their bonus to allocate across one spending account and four separate savings accounts: three accounts earmarked with their chosen goals and one account for "all other savings" (see SM for details on experimental stimuli). As the key dependent variable, we calculated the percentage of participants' bonus that they allocated towards savings.

Results and discussion

We found that setting savings goals – and thus drawing attention to multiple reasons for saving – had no effect on subsequent allocations to savings, whereas budget partitioning had a large effect. Participants in the 'goals' condition allocated 53.6% of their bonus to savings, on average (*SD* = 28.1%), and participants in the control condition allocated 52.7%, on average (*SD* = 31.2%). This difference was not statistically significant ($\text{difference}_{\text{goals - control}}$ = -0.89, *SE* = 4.79, *p* = .85, 95%CI[-10.32, 8.55]). The 'goals+partitioning' condition led participants to allocate 71.3% of their bonus to savings, on average (*SD* = 31.2%). The difference between the 'goals+partitioning' condition and 'goal' condition was statistically significant ($\text{difference}_{\text{goals+partitioning - goals}}$ = 17.68, *SE* = 4.61, *p* < .001, 95%CI[8.59, 26.76]) as was the difference between 'goals+partitioning' condition and the control condition ($\text{difference}_{\text{goals+partitioning - control}}$ = 17.68, *SE* = 4.61, *p* < .001, 95%CI[8.59, 26.76]). See Figure 4.

Figure 4. Study 4: effects of condition on percentage of average monthly income allocated to savings

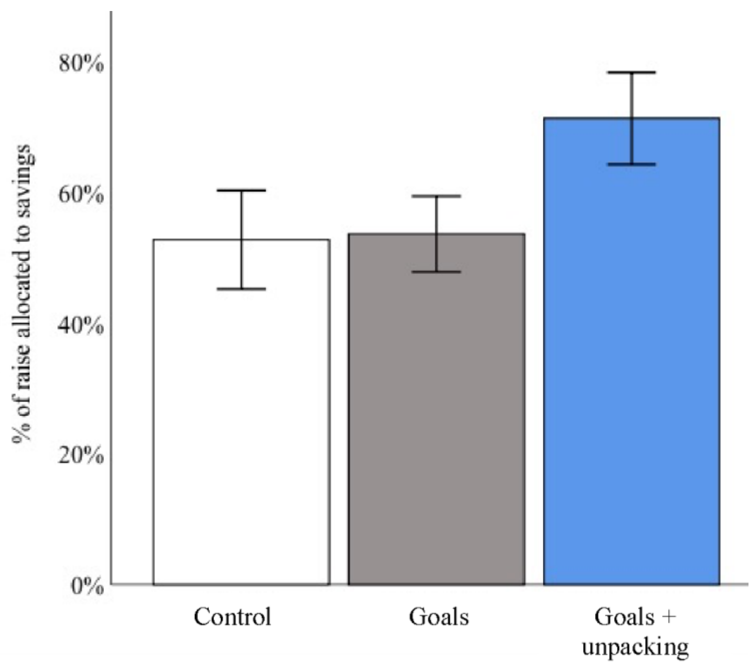

Notes. Displaying means and 95% confidence intervals.

These results indicate that the observed increase in allocation to savings is driven by budget partitioning and not by the act of setting savings goals or by drawing attention to multiple reasons for saving. Goal-setting was held constant in the 'goals' and 'goals+partitioning' conditions (in both conditions participants were asked to take a moment to think about their reasons for savings and select 3 savings goals from the list of 12), yet we only observed an increase in savings when participants explicitly allocated money into partitioned savings sub-categories. These results are consistent with participants using a 1/n heuristic, suggesting that budget partitioning influenced savings by changing the process by which people make allocation decisions.

Furthermore, these results indicate that prompting people to set savings goals is not sufficient to influence their subsequent financial allocation decisions. This does not necessarily mean that goal-setting is inconsequential. Rather, the effects of setting savings goals may only be realized when these goals are built into decision-makers' subsequent choice architecture. We address the implications for the design of financial planning and budgeting software tools in more detail in the General Discussion.

STUDY 4: the effects of additional savings categories

In Study 4, we further explore the budgeting decision process by varying the number of savings categories. If people are, in fact, relying on a 1/n heuristic then partitioning savings into a greater number of saving sub-categories should increase total allocations to savings. In this study, we present participants with a budgeting tool that is partitioned into 2, 4, or 11 savings sub-categories.

Method

Study 4 was a preregistered experiment conducted with 316 online participants (M_{age} = 34.84, SD = 10.06; 37% women; median annual income = $47,000; 93% employed full-time). After entering their annual income for the previous year, participants were randomly assigned to select 1, 3, or 10 savings goals from the same list of 12 savings goals used in the previous study. Participants were then asked to imagine they received a 20% raise (after tax) on their annual income. We calculated each participant's raise amount based on their reported income, and then asked them to decide how much of their raise they wanted to allocate to savings. In the ‘2-savings-categories’ condition, participants allocated this raise across 3 budget categories: their chosen savings goal, ‘all other savings,’ and ‘all spending.’ In the ‘4-savings-categories’ condition, participants allocated their raise across 5 budget categories: their 3 chosen savings goals, ‘all other savings,’ and ‘all spending.’ In the ‘11-savings-categories’ condition, participants allocated their raise across 12 budget categories: their 10 chosen savings goals, ‘all other savings,’ and ‘all spending.’ As the key dependent variable, we calculated the percentage of the hypothetical raise participants allocated to saving.

Results and discussion

The results of a preregistered ANOVA show that partitioning savings into a greater number of categories increased total allocations to saving, $F(2, 313) = 32.24$, $p < .001$. Participants allocated the largest percentage of their raise to savings in the 11-savings-goals condition ($M = 90.42\%$, $SD = 16.43\%$), followed by the 4-savings-goal condition ($M = 84.54\%$, $SD = 12.90\%$), and then by the 2-savings-goals condition ($M = 73.97\%$, $SD = 16.05\%$). All of the pairwise comparisons showed statistically significant differences (11-savings-goals versus 4-savings-goals: $F(1, 313) = 7.67$, $p = .006$; 4-savings-goals versus 2-savings-goals: $F(1, 313) = 24.46$, $p < .001$; 11-savings-goals versus 2-savings-goals: $F(1, 313) = 63.08$, $p < .001$). See Figure 5.

Figure 5. Study 4: effects of condition on percentage of raise allocated to savings

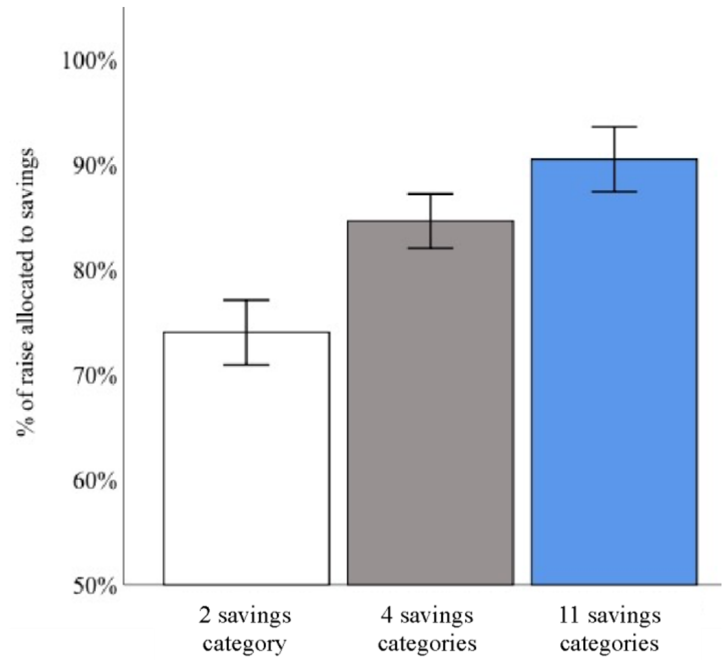

Notes. Displaying means and 95% confidence intervals.

The results of Study 4 provide further evidence that people rely on a 1/n heuristic when allocating money under these budgeting conditions. As savings is partitioned into a greater number of categories, participants' over-reliance on a 1/n heuristic leads to greater overall savings allocations. While there may be diminishing effects – or perhaps even a reversal of the effect – beyond a certain number of savings categories, these results align with an indiscriminate application of a 1/n heuristic.

Thus far, we have examined whether the partitioning of budgeting tools can influence financial allocation decisions. Our results indicate that partitioning may be a particularly effective 'nudge' that can be used by product designers and marketers to influence how people allocate money. However, a partitioning intervention will not be effective in all circumstances. We predict that people will only rely on a 1/n heuristic under specific conditions. That is, only when the choice architecture of a given budget tools prompts people to consider their whole budget. Typically, people do not think about their whole budget when making financial allocation decisions. Rather, they focus on isolated trade-offs among a small subset of options (e.g., “Would I rather spend money on a vacation or contribute to my savings?”). This narrow frame is how people approach most resource allocation decisions. People often only pay attention to information that is explicitly presented in the choice architecture or immediately salient (e.g., this is related to the concept: “what you see is all there is”; Kahneman, 2011; Kahneman, Lavallo, & Sibony, 2011). Similarly, when people consider allocating money towards any given purpose, they often only think about related alternatives or the “next best use,” neglecting outside options and failing to consider the full range of opportunity costs (Slovic, 1972; Kahneman & Fredrick, 2002; Okada & Hoch, 2004; Frederick et al., 2009). When adopting this kind of narrow frame, the 1/n heuristic is not useful as a simplifying strategy. Therefore, we predicted that people would only use a 1/n heuristic when they adopt a broader frame. That is, when they consider their whole budget and need a shortcut to reduce the complexity of their allocation decision. In the Studies 5-7 we identify three features of household budgets that prompt people to consider their whole budget and therefore be more likely to rely on a 1/n heuristic. These three features are: 1) fixed sum allocation, 2) exhaustive budget categories, and 3) simultaneous allocation. We examine each of these features and how they can be implemented into budgeting tools.

STUDY 5: Fixed sum budget allocation

In Study 5, we examined whether people would be more likely to rely on a 1/n heuristic in household budgeting decisions when allocated an explicit, fixed sum of money. When people decide how much money to contribute to their savings, sometimes they are dividing an explicit, fixed sum (e.g., a bonus, stimulus check, sales commission, or their monthly income) and other times they do not have any specific fixed sum in mind (e.g., thinking about their general financial circumstances and how much money they feel like they can afford to save right now). We predicted that people would be more likely to use a 1/n heuristic if a relevant fixed sum, such as their monthly income, was more salient at the time of the budget allocation decision.

Method

Study 5 was a preregistered experiment with 948 online participants in the United States, recruited from Prolific (M_{age} = 40.49, SD = 12.34; 51% women; median monthly income = $3,475; 81% employed full-time). In order to be eligible for this study, participants had to report that “after [they] pay their bills, [they] typically have some money left over,” and that they would consider “signing up for an online service that helps [them] save a part of [their] monthly income.” Participants who met these inclusion criteria were given the following instructions: “Imagine you just signed up for a financial service that can help you grow your savings. This service allows you to set up an auto-deposit towards your savings goals … the following questions will ask about how you would use this service.” Next, participants to selected 3 savings goals from a list of 12 (using the same list as in the previous two studies) and then decided how much money to automatically deposit towards these goals each month. Participants were randomly assigned to one of six conditions to decide on their monthly auto-deposit in a 2 (savings partition: 1 vs 3 savings accounts) x 3 (income salience: none vs low salience vs high salience) experimental design. The first treatment factor varied whether participants allocated money towards a combined savings account (with their 3 chosen savings goals listed in brackets) versus separate accounts for each savings goal (see SM for experimental stimuli). The second treatment factor varied the salience of participants’ monthly income at the time of this allocation decision. This treatment factor had three levels: 1) no mention of monthly income, 2) participants were asked to merely think about their average monthly income in advance and at the time of the allocation decision, 3) participants were asked to report their average monthly in advance and this number was presented back to them at the time of allocation decision. See SM for more detail on the experimental stimuli. The key dependent variable was the combined amount of money participants' chose to auto-deposit towards their three chosen goals each month.

Results and discussion

The results support our prediction that partitioning effects are stronger when people allocate an explicit, fixed sum of money. We found a significant main effect of partitioning savings (one-way ANOVA: $F(1, 933) = 18.88$, $p < .001$) and a significant main effect of income salience (one-way ANOVA: $F(2, 933) = 12.89$, $p < .001$). We also found our predicted result in our primary comparison of interest (conditions 4 versus 6): partitioning savings with no mention of monthly income ($M = 13.3\%$, $SD = 13.0\%$) versus partitioning savings while prompting participants to enter their monthly income ($M = 20.6\%$, $SD = 20.6\%$). We found that people allocated an extra 7.3% of their monthly income to savings when the budgeting tool included this explicit reference amount, $F(1, 310) = 13.37$, $p < .001$. These results indicate that people are more likely to use a 1/n of heuristic when they have a fixed reference sum in mind.

Figure 6. Study 5: effects of condition on percentage of monthly income allocated towards a monthly savings auto-deposit

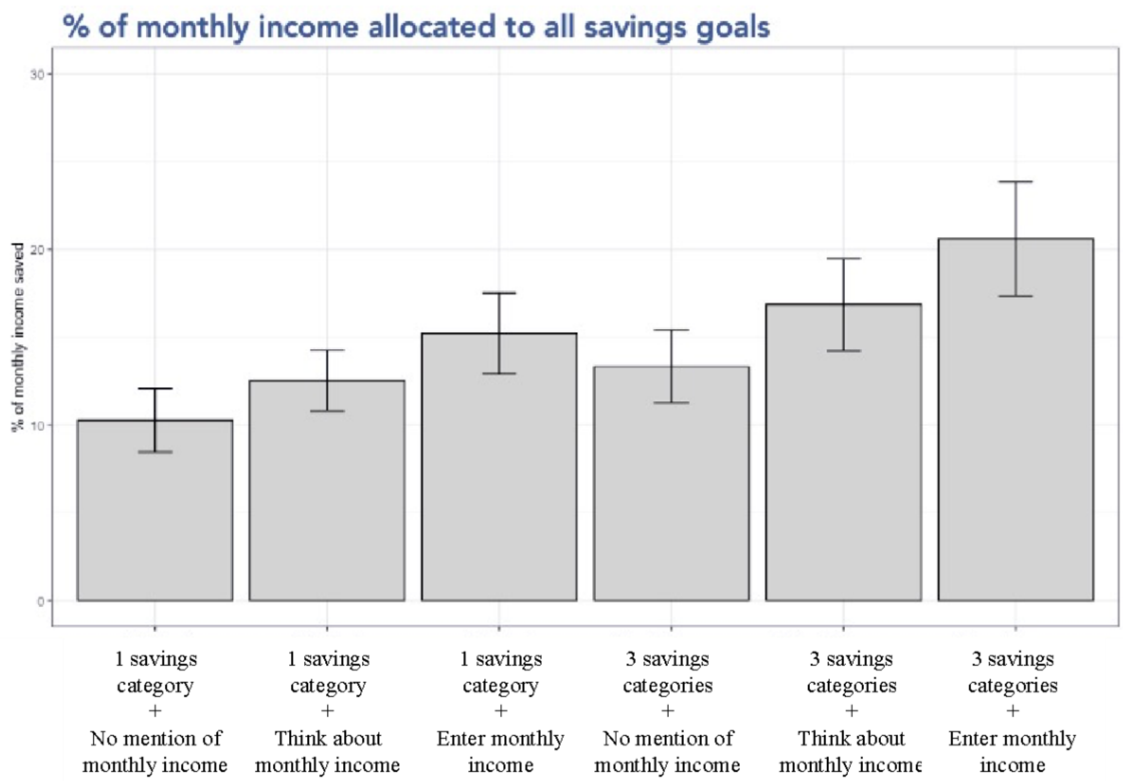

Notes. Displaying means and 95% confidence intervals.

STUDY 6: Allocating money using exhaustive versus non-exhaustive budget categories

In Study 6, we examined whether people were more likely to rely on a 1/n heuristic when they allocated money across an exhaustive set of budget categories. Often, budgeting decisions are made using a non-exhaustive set of categories. For example, people may decide how much of their monthly income to allocate towards saving and leave the remainder in their checking account to be used for bills and all other spending. This is commonplace as people often receive monthly paychecks as direct deposits into their checking account and then subtract from this account to allocate money towards specific purposes.

We predicted that when people allocated a fixed sum of money over an exhaustive set of categories with no default, they will be more likely to use a 1/n heuristic and exhibit bias towards even allocation across identified budget categories.

Method

Study 6 was a preregistered experiment with 681 participants in the United States recruited from Amazon's Mechanical Turk (M_{age} = 40.0, *SD* = 13.0; 50.3% women; median monthly income = $3000; 78.4% employed full-time). Participants first entered their average monthly income. Then, they were asked to imagine they were setting up a direct deposit at a new bank and they needed to decide how to allocate their monthly income across their new savings and spending accounts. We randomly assigned participants to one of four conditions in a 2 (saving accounts: 1 versus 4) x 2 (budget categories: non-exhaustive versus exhaustive) experimental design.

The first treatment factor varied whether participants allocated their monthly income into a single general savings account versus 4 savings accounts, earmarked with common savings goals: saving for emergencies, saving for upcoming expenses and purchases, saving for retirement, and all other savings). In order to hold information about reasons for saving constant across conditions, participants who allocated money into the single savings account were informed that "this includes saving for emergencies, upcoming expenses and purchases, retirement, and all other savings."

The second treatment factor varied whether participants allocated money into spending account versus the having the spending account as the default. In the conditions where the spending account was the default, we informed participants that "any amount not transferred into savings will remain in your spending account." So, when the spending account was presented as an explicit budget category, the allocation task involved an exhaustive set of options and participants had to allocate their entire monthly income, and when the spending account was absent, the allocation task involved a non-exhaustive set of options, such that participants decided how much money to deduct from their spending account towards savings.

Results and discussion

The results of a preregistered two-way ANOVA show a significant main effect of partitioning savings, $F(1, 677) = 16.86$, $p < .001$, and a significant main effect of including an exhaustive set of budget categories, $F(1, 677) = 25.71$, $p < .001$. We did not find a significant interaction effect, $F(1, 677) = .99$, $p = .319$.

Our key comparison of interest was between conditions 3 and 4. Here, we observe a significant difference such that participants allocated 23.5% (SD = 24.5%) of their monthly income towards saving when the budgeting tool was partitioned with a non-exhaustive set of budget categories and 34.6% (SD = 27.1%) with exhaustive set of categories, $F(1, 310) = 13.84$, $p < .001$.

Figure 7. Study 6: effects of condition on percentage of monthly income allocated to savings

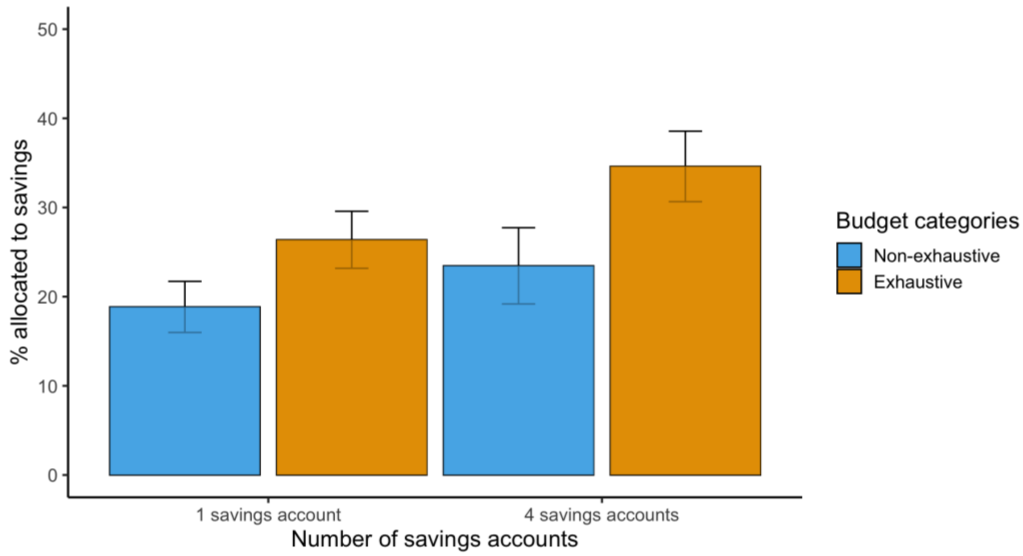

Notes. Displaying means and 95% confidence intervals.

STUDY 7: The effects of budgeting sequentially versus simultaneously

In Study 7, we examined another feature of household budgeting that might lead people to rely on a 1/n heuristic. We predicted that people would be more reliant on a 1/n heuristic when they follow a *simultaneous,* versus sequential, budgeting procedure. Often, people make budgeting choices *sequentially*, for instance, by first allocating money to their primary financial obligations (e.g., non-discretionary expenses like rent, heating, and car payments), then considering secondary expenses (e.g. food and groceries), and finally considering discretionary or hedonic purchases (e.g. clothing, jewelry, entertainment). We predicted that people would be less likely to rely on a 1/n heuristic when following a sequential budgeting procedure because this heuristic is not adaptive under these conditions (Shah & Oppenheimer, 2008). That is, it does not effectively reduce cognitive effort when making decisions in a sequential budgeting procedure because it requires people hold multiple sums in working memory from one decision to the next. However, a 1/n heuristic can be an effective way to reduce cognitive effort when people follow a *simultaneous* budgeting procedure in which they spread a fixed sum of money across budget categories all at once (e.g., on a single screen in a budgeting tool). Therefore, we predicted that people would be more likely to rely on a 1/n heuristic, and therefore display more partition-dependent preferences, when they used a budgeting tool that presented all budget categories together and prompted people to allocate across these categories at one time.

Method

Study 7 was a preregistered experiment conducted with 312 online participants (M_{age} = 33.64, *SD* = 9.35, 38% women; $M_{monthly\ income}$ = $5,231, *SD* = $6,182; 85% employed full-time). In order to be eligible, participants had to indicate that they had some discretionary money to save. Participants reported their average monthly income and then selected three savings goals from the same list of 12 goals that was presented in Studies 2-5. Then, we asked participants to construct a realistic household budget for the upcoming months. To do so, participants allocated their average monthly income across four categories: three savings categories labeled with their chosen savings goals and one spending category labeled as “all other uses, including bills, debt repayments, and all spending.” Half of participants allocated their monthly income sequentially, with each of the four categories presented on a separate page. If participants were not satisfied with the final budget after completing the sequential allocation, they were able to repeat the process and make changes. The other half of participants allocated their monthly income simultaneously, with all budget categories presented on a single page. In both conditions, the budgeting procedure was explained to participants ahead of time. See Figure 8 for a study diagram.

Figure 8. Study 7: diagram of sequential and simultaneous budgeting procedures

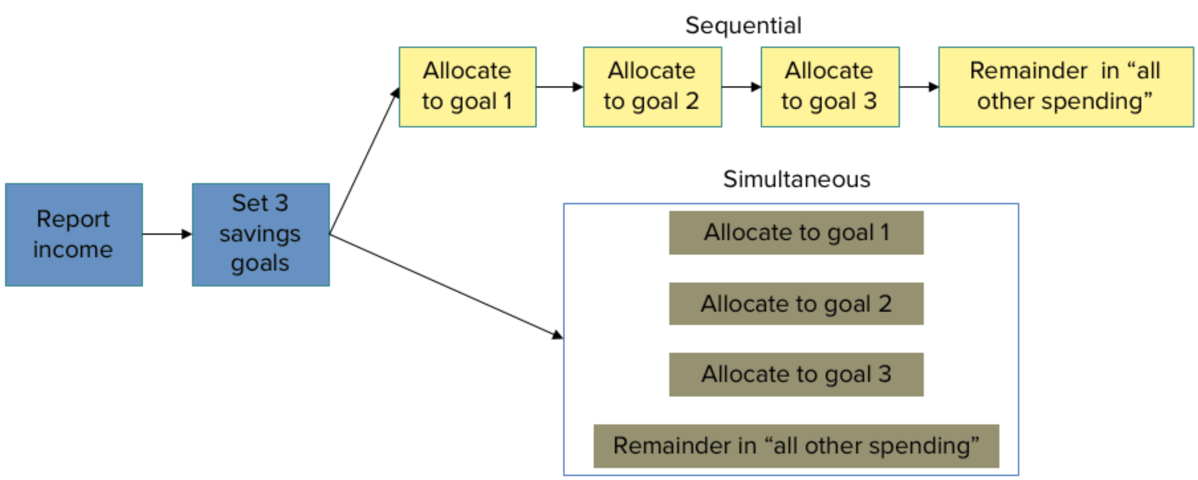

Notes. The 3 specific savings goals were selected by participants from a list of 12 common goals. Order of spending and savings categories was counterbalanced.

Results and discussion

In a preregistered ANOVA, we found that participants in the ‘simultaneous’ condition allocated an additional 16% of their monthly income to savings, relative to participants in the 'sequential' condition, ($M_{simultaneous}$ = 43.47%, SD = 29.55%; $M_{sequential}$ = 27.63%, SD = 22.74%; $F(1,311) = 27.68$, $p < .001$). See Figure 9. As in the previous studies, we conducted preregistered robustness checks, controlling for income(log), age, and education. We find that the results are substantively unchanged after including these controls (see SM). We also observe no interaction effect between condition and income with respect to effects on savings allocation, $b = 9.946$, $t(3,311) = 1.046$, $p = .296$, CI(b) = [-8.757, 28.650].

Figure 9. Study 7: effects of condition on percentage of average monthly income allocated to savings

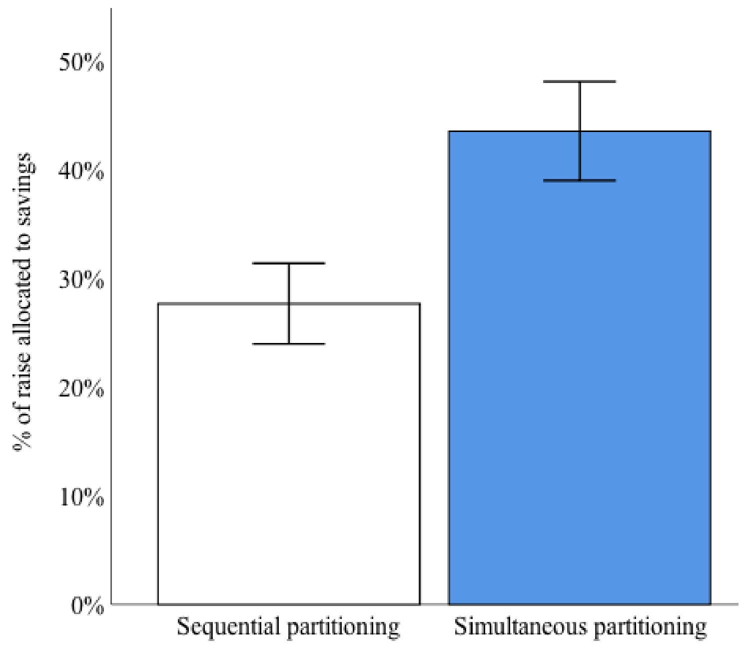

Notes. Displaying means and 95% confidence intervals by condition.

These results suggest that presenting budget categories simultaneously may increase the likelihood that people rely on a 1/n heuristic. Here, the number of budget categories was held constant such that all participants allocated their monthly income across 4 categories. The only difference between conditions was whether participants allocated money sequentially or simultaneously. Sequential allocation is not amenable to 1/n of heuristic since this would require participants to either calculate a running tally of total allocations from one page to the next, or plan all of their decisions in advance. It is more likely that, for each savings goal, participants asked themselves: “how much money am I able to allocate to this savings goal right now? In contrast, the simultaneous condition enabled participants to effectively reduce cognitive effort using a 1/n heuristic.

STUDY 8: Combining each feature of budgeting

In Study 8, we combined the insights from previous studies to develop and test a savings nudge based on budget partitioning. We conducted this study in November when many U.S. workers were anticipating a holiday or year-end bonus. We asked participants to allocate their upcoming bonus into their household budget using one of five budgeting tools.

Based on the results of the previous studies, we preregistered three predictions. First, partitioning savings accounts into multiple sub-accounts will increase total allocation to savings relative to a single savings account and relative to setting multiple savings goals. Second, people will save more when allocating money across multiple savings accounts simultaneously, as opposed to sequentially. Third, people will allocate more money to savings when savings sub-accounts are presented alongside an overall spending account such that people are allocating a fixed sum of money across an exhaustive set of options.

Method

Study 8 was a preregistered experiment conducted with 1,022 online participants (M_{age} = 36.19, *SD* = 12.33, 47% women; median annual = $45,000; 71% employed full-time). In this study participants were asked to allocate a hypothetical 20% year-end bonus into their household budget. Participants indicated their annual income, then we told them the dollar amount of their holiday bonus. We assigned participants to one of five procedures to allocate this windfall into their household budget. In the control condition (1), participants simply decided how much of this bonus to allocate to savings using a single, overall savings category. In the 'savings goals' condition (2), participants chose 4 savings goals from the same list of 12 common saving goals, and then decided how much to allocate into the same overall savings category. In the 'sequential' condition (3), participants again chose 4 savings goals and then allocated money across each goal plus an 'all other savings' category in a sequential process (i.e., participants allocated money to each of their goals one-by-one on separate screens). The 'simultaneous' condition (4), was identical to the previous condition except that participants allocated money to all of their chosen savings goals simultaneous (i.e., on a single screen). Finally, the 'spending account' condition (5), was identical to the 'simultaneous' condition, except that we added an all-purpose spending category to the budgeting tool so that participants allocated the entire bonus across 4 savings categories and 1 spending category all on one screen. See SM for diagrams of each budgeting procedure.

Results and discussion

In a preregistered ANOVA, we found a significant omnibus effect of condition on savings allocation, $F(4, 1017)=42.96$, $p < .001$. In planned pairwise comparisons, we found that setting savings goals had no impact on savings allocations whereas partitioning had a large effect, especially simultaneous partitioning (See Figure 14 and Table 10). In preregistered regressions we found that these effects hold including a set of financial controls (income, subjective financial wellbeing, and financial literacy score) as well as a set of demographic controls (age, gender, and education). See SM for regression results.

Figure 10. Study 8: effects of condition on percentage of bonus allocated to savings

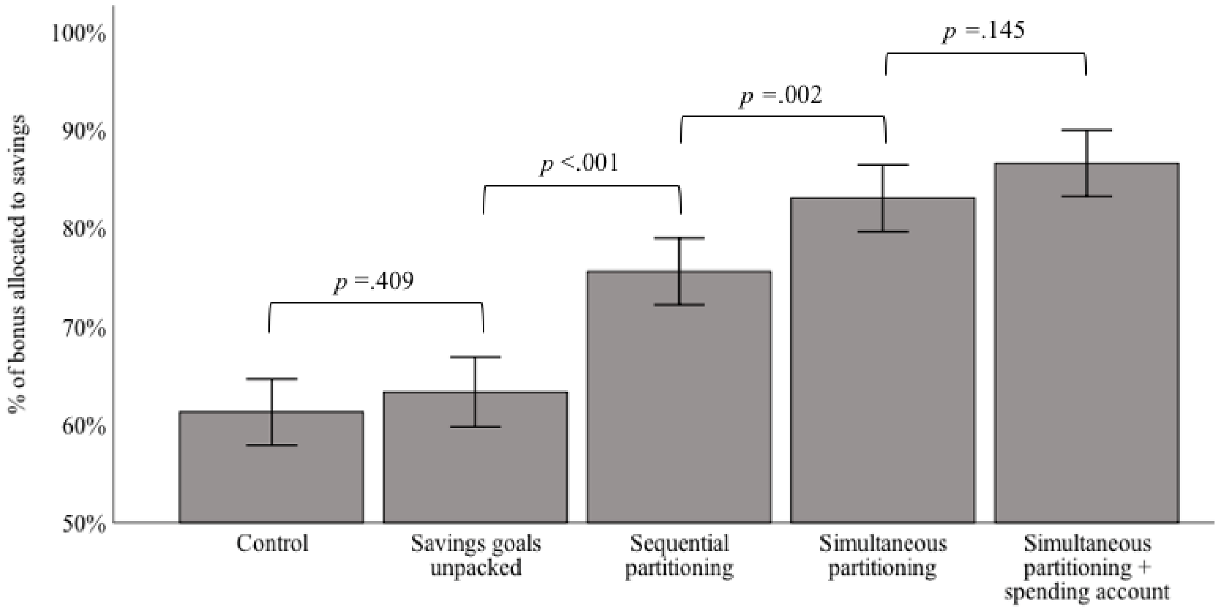

Note. Displaying means and 95% confidence intervals. P-values correspond to preregistered planned contrasts in a general linear model with LSD corrections for multiple comparisons.

Table 1. Study 8: all pairwise comparisons

Pairwise comparison		Mean difference	*p*	95% CI
Savings goals unpacked	Control	2.06	.409	[-2.83, 6.95]
Sequential partitioning	Control	14.34	<.001	[9.56, 19.12]
Simultaneous partitioning	Control	21.78	< .001	[17.00, 26.56]
Simultaneous+spending	Control	25.34	< .001	[20.57, 30.11]
Sequential partitioning	Savings goals unpacked	12.28	< .001	[7.38, 17.19]
Simultaneous partitioning	Savings goals unpacked	19.73	< .001	[14.82, 24.63]
Simultaneous+spending	Savings goals unpacked	23.28	< .001	[18.38, 28.18]
Simultaneous partitioning	Sequential partitioning	7.44	.002	[2.64, 12.24]
Simultaneous+spending	Sequential partitioning	10.99	< .001	[6.21, 15.78]
Simultaneous+spending	Simultaneous partitioning	3.56	.145	[-1.23, 8.34]

Note. Reporting mean differences (measured in percentage points) for each pairwise comparison. P-values and 95% confidence intervals are calculated using an LSD correction for multiple comparisons.

These results support two of our three preregistered predictions. Aligning with the results of Study 3, we found that partitioning savings into multiple sub-categories led to a significant increase in total saving, whereas setting savings goals without partitioning had no effect.

Aligning with the results of Study 7, simultaneous partitioning led to more saving than sequential partitioning. However, we did not find evidence in support of our third prediction. We expected that adding a spending account – and therefore providing an exhaustive set of budget categories whereby participants allocated the entire bonus – would increase the total allocation to savings, consistent with the results of Study 6. However, we find no significant increase in savings in condition 5 relative to condition 4 (mean difference = 3.56 [-1.23, 8.34]. It is possible that we did not observe a significant effect of due to a ceiling effect. Participants using a budgeting tool with 4 savings categories and simultaneous allocation (condition 4) already allocated XX% of their bonus to savings, on average, and therefore it may be difficult for any features of budgeting choice architecture to further increase savings.

GENERAL DISCUSSION

This research shows that the choice architecture of digital budgeting tools can play a critical role in shaping financial allocation decisions. Across eight experiments, we identified specific design features that cause some budgeting tools to encourage more savings and others to encourage more spending. Studies 1 and 2 found that partitioning savings into multiple sub-categories can drastically increase the total amount of money people allocate towards savings, relative to a single all-purpose savings category. Study 3 isolated the effects of budget partitioning from financial goal-setting and found that setting savings goals had no effect on subsequential financial decisions unless goals were partitioned in during the allocation decision. Study 4 showed that partitioning savings into a greater number of sub-categories in a budgeting tool led to more overall savings, consistent with a 1/n heuristic. Studies 5-7 found that the effects of partitioning were stronger in budgeting tools that included a relevant reference amount (e.g., one's monthly income), an exhaustive set of categories, and a simultaneous allocation procedure

– all features that prompt people to consider their whole budget, rather than isolated trade-offs. Finally, Study 8 showed that combining these features can further increase partitioning effects, and therefore these features can be used to design an effective, non-coercive savings nudge.

Across these eight experiments, we observe effect sizes of budget partitioning ranging from medium to large (d = .55 to d = 1.58). We would not expect to find such large effects in a field context. Indeed, we observed the smallest effect size in Study 2, when participants made an incentive compatible choice about how to allocate a $1000 cash prize. However, in this study, we still found that the partitioning manipulation led participants to allocate an extra $159 to saving, representing a medium to large effect size (d = .55). Importantly, all of these partitioning interventions produced reliable effects merely by changing the elicitation of savings decisions while preserving participants' free choice and holding constant the information provided. Based on these results, partitioning appears to be more consequential than goal-setting, which is the basis of most online banking and financial software applications, and more consequential than the well-established nudges that have been documented in the financial decision-making literature. Even highly successful nudges in this domain typically have small effect sizes. The nudges that do yield larger effect sizes usually required automating decisions and, therefore, they are not entirely choice-preserving (for more detail see: Beshears & Kosowsky, 2020).

We theorize that the partitioning effects observed in our studies are caused by participants relying on a 1/n heuristic, and the pattern of results we observe across eight experiments is consistent with this theory. However, one limitation of our experiments is that we do not directly measure people's decision process. We did not employ process-tracing methods such as 'thinking aloud' protocols, eye-tracking, and mouse-tracking because it remains unclear how these methods could feasibly measure a 1/n decision process. Self-reports of one's own decision processes are usually unreliable and using a 1/n heuristic is not necessarily correlated with eye movements or fixation times, which is why this research has focused on inferential experimental approaches. Nevertheless, future research could build on this work by collecting novel process-tracing data to determine the conditions under which people rely on a 1/n heuristic in resource allocation decisions.

While we do not have process evidence for a 1/n heuristic, we do present robust evidence for a partition-dependence bias in household financial allocation decisions. We can reliably predict this bias in preregistered experimental protocols, and we identify three specific features of the choice architecture that are critical moderators of this bias. These insights have implications for the design of fintech software and may be generalizable to other resource allocation decisions. For example, many financial software applications involve setting goals, creating separate categories or accounts, and dividing money. Most major banks provide an online banking platform that allows users to create multiple earmarked savings accounts and set up auto-deposits to divide incoming money across these accounts. There has also been a proliferation of alternative consumer financial technology products that provide users with automated advice, investment options, trading platforms, and budgeting software. Many of these digital financial tools may be inadvertently influencing users' allocation decisions because of the ways in which they partition people's finances. Indeed, this research highlights that there is no "neutral" way to partition financial categories nor any "neutral" allocation procedure. Product designers should be aware of how these design choices influence users and, where possible, they should aim to align the choice architecture with users' own stated goals. For instance, if users have a goal to save more money, online banking, investment, and budgeting tools could encourage these users to partition their finances such that they allocate money into a greater number of savings sub-categories using a relevant reference amount, an exhaustive set of categories, and a simultaneous allocation procedure.

Overall, this research adds to our understanding of how people make financial allocation decisions and the ways in which financial tools can shape these decisions. The partitioning of the budgeting tools that we as consumers use to manage our finances may play a critical role in how we allocate money and, ultimately, our long-term financial wellbeing.