

Spending Too Little in Hard Times

Alessandro Del Ponte and Peter DeScioli
Stony Brook University

Abstract

People's decisions to consume and save resources are critical to their wellbeing. Previous experiments find that people typically spend too much because of how they discount the future. We propose that people's motive to preserve their savings can instead cause them to spend too little in hard times. We design an economic game in which participants can store resources for the future to survive in a harsh environment. A player's income is uncertain and consumption yields diminishing returns within each day, creating tradeoffs between spending and saving. We compare participants' decisions to a heuristic that performed best in simulations. We find that participants spent too much after windfalls in income, consistent with previous research, but they also spent too little after downturns, supporting the resource preservation hypothesis. In Experiment 2, we find that by varying the income stream, the downturn effect can be isolated from the windfall effect. In Experiments 3-4, we find the same downturn effect in games with financial and political themes.

JEL Codes: **C63** (Computational Techniques – Simulation Modeling); **C91** (Design of Experiments – Laboratory; Individual Behavior); **D15** (Intertemporal Household Choice - Life Cycle Models and Saving); **D90** (Micro-Based Behavioral Economics); **E21** (Consumption; Saving)

Keywords: Spending; saving; ownership; mental accounts; intertemporal choice; ecological rationality

Introduction

People's decisions to consume and save resources are critical to their prosperity, health, and even survival when scarcity is severe. In diverse cultures throughout history, people have stored food and supplies to sustain them through difficult times. More than 10,000 years ago, Neolithic foragers stored wild barley in large well-designed granaries, and these storage technologies fostered the major transition from foraging to agriculture (Kuijt & Finlayson, 2009). In modern economies, people store wealth in bank accounts and financial assets, and they draw on these funds to manage hardships such as unemployment, medical emergencies, or natural disasters. Here we investigate the psychological mechanisms that help or hinder people as they seek to efficiently manage resources.

Previous research in behavioral economics emphasizes a major cognitive limitation: Experiments show that people's decisions are myopic and focus on the present while excessively discounting the future, hence they spend too much and save too little (Berns et al., 2007; Frederick et al., 2002). This finding resonates with the fact that many people do not appear to save enough money. For instance, one report found that roughly half of U.S. consumers did not have sufficient savings to cover their expenses for three months if their income was interrupted (Larrimore, Dodini, & Thomas, 2016). However, in some cases, people are able to save substantial amounts of money, food, and other vital resources. Moreover, these saving behaviors are found widely across cultures and they play a key role in economic prosperity (Smith, 1776/1904). Hence, the prevalence of saving raises a question about which psychological mechanisms oppose people's tendency to discount the future and thus help them save resources in the present.

We propose an additional mechanism behind saving decisions that is based on the psychology of resources rather than time. When someone saves a resource such as money or food, they gain access, control, and ownership of it, whereas spending a resource means losing it. In order to maintain access to valuable resources, an individual should accumulate and store resources that could become scarce and then consume these reserves sparingly through difficult times. Moreover, people's motive to preserve their resources could potentially cause them to spend too little, particularly during a period of scarcity that is actually transitory rather than indicative of a downward trend in the availability of resources.

The psychology of resource storage is likely to be closely connected to people's sense of ownership. Namely, once resources are stored, ownership motivates people to defend them against others. Previous research argues that people have an evolved sense of ownership that motivates them to fight to defend what they own (Boyer, 2015; DeScioli & Wilson, 2011; Maynard Smith 1982; Sherratt & Mesterton-Gibbons, 2015; Stake, 2004), to attribute greater value to owned objects (Jones & Brosnan, 2008; Kahneman et al., 1990; Morewedge & Giblin, 2015), and to bargain more aggressively for profits that they earned through effort or skill (Cherry et al., 2002; Hoffman et al., 1994). We suggest that the psychology behind storing resources is a precursor and component of the ownership strategies that people use to manage conflicts over valuable goods. Hence, similar to owned objects, once someone stores a resource, they might value it beyond equivalent resources that are not yet or only recently acquired. Moreover, this preservation motive might be enhanced by cues of scarcity, such as a sudden decrease in the availability of a resource. On the other hand, cues of abundance are expected to reduce the preservation motive, leading to greater spending.

The motive to conserve stored resources could counteract and possibly even overpower the tendency to discount the future. For instance, one previous study found that participants sometimes prefer to pay for products sooner rather than later—essentially discounting the present—in order to end the aversive experience of parting with their money (Prelec & Loewenstein, 1998). In fact, people’s greater value for owned resources could even cause them to spend too little in some cases. Namely, when hardship strikes, this may act as a cue for a period of scarcity, which could make someone too reluctant to deplete their reserves and thus spend too little to efficiently smooth consumption. Moreover, it could be ecologically rational (Todd & Gigerenzer, 2007) to tighten the budget in hard times because a period of scarcity could be predictive of future scarcity in many real-world environments; hence, people might intuitively infer an enduring decrease in income (Gallistel et al., 2014) even if the decrease is actually transitory.

Mental accounting and the psychology of resource preservation

The theory of mental accounting holds that people mentally separate their wealth into different accounts, and they spend money differently depending on which account it is contained in (Thaler, 1985, 1999). Specifically, people have different mental accounts for new income, current wealth, and future income, along with various other accounts that are earmarked for special purposes such as buying a car or house. A key point of the theory is that people’s mental accounts violate the economic principle of fungibility, which means that the way someone labels wealth is irrelevant for how it can be most efficiently spent.

Research on the psychology of resources and ownership can help deepen the theory of mental accounting. One of the main ideas from mental accounting is that people are more willing

to spend money from their recent income than their accumulated wealth (Thaler, 1999). The psychology of resources can explain why people make this distinction, namely because they have stronger motives for preservation for their accumulated wealth than for new income.

A growing literature in psychology examines the cues that shape people's judgments about who owns what, including the strength of ownership and how people weigh and reconcile conflicting cues that favor different potential owners (Boyer, 2015; DeScioli & Karpoff, 2015; DeScioli, De Freitas, & Karpoff, 2017). Much of this research studies participants' judgments about ownership dilemmas and examines how ownership is conferred by key events such as such as finding, creating, purchasing, earning, maintaining, and giving (Blake & Harris, 2009; Blumenthal, 2010; Friedman & Neary, 2008; Kanngiesser, Gjersoe, & Hood, 2010; Kanngiesser & Hood, 2014; Kim & Kalish, 2009). Particularly relevant for mental accounts, one key factor is the time of possession: People attribute less ownership for a recently acquired object than an object held for a longer span of time (DeScioli & Wilson, 2011; Stake, 2004). This might explain why people are more willing to spend new income than their accumulated wealth.

Meanwhile, evolutionary psychology illuminates why humans have a sense of ownership in the first place (Boyer, 2015; DeScioli & Wilson, 2011; Stake, 2004). Specifically, the psychological systems surrounding ownership evolved to motivate people to accumulate and manage reserves of valuable resources and to guard them against intruders. Many other animal species show analogous behaviors including caching resources, defending them, and using conventions such as first possession to resolve disputes over resources (Brosnan, 2011; Kokko, Lopez-Sepulcre, & Morrell, 2006; Maynard Smith, 1982; Sherratt & Mesterton-Gibbons, 2015; Vander Wall, 1990).

Hence, the theory of mental accounting can be integrated with the psychology and evolutionary biology of resource management and ownership. People's motives to stockpile and defend resources explain how in some cases they can overcome their myopia and save some resources for the future. In fact, people's motive to preserve their savings might even sometimes cause them to spend too little in hard times, when consumption requires depleting and losing their reserves.

Consumer spending and excessive sensitivity to transitory income shocks

The idea that people could spend too little in hard times has some support in a large literature in economics on consumer spending. This literature finds that a typical consumer's spending is excessively sensitive to both increases and decreases in recent income, even when these shocks are predictable and transitory, deviating from economic theories of rational consumption (reviewed in Jappelli & Pistaferri, 2010). For instance, consumers spend more money after predictable gains in income such as from tax rebates, and they spend less money after predictable decreases such as reduced income due to retirement (Gelman et al., 2014; Parker et al., 2013; Poterba, 1988; Souleles, 1999; Wilcox, 1990; Wilde & Ranney, 2000). In one recent study, researchers analyzed monthly expenditures from bank account information for 200,000 U.S. households in which someone lost their job and received unemployment insurance (Ganong & Noel, 2017). After job loss, a typical consumer spent 6% less on nondurables like groceries and medicine; six months later, when they predictably exhausted the unemployment benefits, their spending steeply declined again by an additional 13%. In another recent study, researchers examined thousands of federal workers whose pay was temporarily reduced by 40% during the 2013 federal government shutdown and then later repaid in full (Baker & Yannelis,

2017); even though the income shocks were transitory, workers reduced their spending by 7% when their pay was cut and then increased spending up again by 12% when they received the larger check with back pay. These and many similar studies show that consumers overreact in both directions to changes in income, even when these shocks are predictable and transitory. However, the interpretation of these patterns is the subject of longstanding debate (Jappelli & Pistaferri, 2010), in part due to the complexity of real-world spending decisions and the limitations of observational methods.

The present experiments

We investigate people's spending decisions in an economic game with two key features. First, players can store resources for future periods of the game. This differs from the predominant method in previous experiments: the intertemporal choice task, in which a participant chooses whether to receive a small reward sooner or a larger reward later (reviewed in Frederick et al., 2002). The intertemporal choice task captures tradeoffs over time but misses another key aspect of saving, the ability to store resources. Some previous economic games have an element of storage (e.g., Ballinger, Palumbo, & Wilcox, 2003; Brown, Chua, & Camerer, 2009; Carbone & Hey, 2004), but the current game goes further by allowing players to find, accumulate, and store resources in a virtual basket. By incorporating storage, we can test whether participants' spending differs for new resources they just found versus resources in storage. The resource preservation hypothesis predicts that participants will be less willing to spend stored resources than newly acquired ones, particularly in times of scarcity.

Second, the game allows researchers to measure participants' performance. The intertemporal choice task does not have a clear performance standard, and researchers often view

these choices as a matter of personal preference. However, performance is a crucial aspect of saving decisions because people's prosperity and wellbeing depend on how efficiently they manage resources. Good saving decisions can make the difference between sufficiency and poverty, shelter and homelessness, health and disease, life and death. In the present game, saving is a matter of survival rather than preference. Performance is based on how many periods a player survives in the game, and a player's choices can be compared to saving heuristics that survive the most periods in computer simulations.

By measuring performance, we can assess whether participants spend too much, too little, or just right in different situations. Previous theories generally predict that people will overspend because they excessively discount the future (Berns et al., 2007; Frederick et al., 2002). Moreover, previous research also finds that people spend even more after windfalls when they receive particularly high income (Arkes et al., 1994; Milkman & Beshears, 2009) due to how they represent unexpected income in their mental accounting (Thaler, 1985, 1999).

The preservation hypothesis makes an additional prediction that people could also spend too little after a downturn in income. An efficient spender will draw enough from their savings in hard times to smooth their consumption over time. However, if people attribute additional value to their stored resources in response to scarcity, then they might sometimes spend less than optimal due to a motive to preserve their savings.

We test these predictions in Experiment 1 by comparing how participants spend new resources versus stored resources. In an economic game, participants receive fluctuating income that they can spend or store, which naturally provides repeated within-subject observations of participants' spending when they have different combinations of new and stored resources. Moreover, we also manipulate between-subjects a key factor in economic theories about saving:

the variance in income. In the ten-zero condition, participants receive an income of 10 or 0 each day (50% chance each), and in the seven-three condition, they receive an income of 7 or 3; hence the average is the same (5 per day) but the variance is greater in the first case. This manipulation varies the magnitudes of windfalls and downturns which is expected to amplify their respective effects on spending. Most relevant here, the high variance environment creates a stronger cue of scarcity because the downturn in income deviates further from the long-term average, and efficient spending requires greater depletion of stored resources.

In Experiment 2, we further vary the stream of income to see if it is possible to isolate the effect of downturns in income from the effect of windfalls. In the game, participants usually receive an income that is close to the long-term average, but this routine is occasionally punctuated by either windfalls or downturns, depending on the condition. Specifically, in a between-subject design, participants in the windfall condition receive an income of 3, 3, 4, or 10 (with equal chances), and participants in the downturn condition receive 6, 7, 7, or 0; hence, both conditions have a long-term average income of 5 but differ in whether the typical income is punctuated by windfalls or downturns, allowing us to observe their effects separately. In particular, the preservation hypothesis predicts greater underspending in the downturn condition, since efficient consumption requires players to deplete more of their savings after steeper downturns.

In Experiments 3 and 4, we change the theme of the game from foraging for food to managing money in a household and a government, respectively. The game parameters and experimental design are otherwise the same as Experiment 2. This allows us to test for windfall and downturn effects when participants' decisions are framed in terms of household finance and the national budget.

Ecological rationality and saving performance

The present experiments assess participants' saving performance relative to a specific goal: surviving more days and earning more money. This approach is more aligned with the concept of ecological rationality (Todd & Gigerenzer, 2007) than with some notions of rationality in behavioral economics that were inherited from neoclassical economics. Specifically, in one stark conception called revealed preference theory (Samuelson, 1948), a rational player can pursue any goal at all (including bankruptcy, starvation, extreme risk-seeking, or death), and the player's rationality requires only consistency. Hence, some strands of behavioral economics have focused specifically on assessing consistency, such as inconsistencies from hyperbolic temporal discounting (Berns et al., 2007). Meanwhile, other economic researchers often use a broader conception of rationality that incorporates specific goals such as profit-maximization, smoothing consumption, Pareto efficiency, or social efficiency.

In contrast to revealed preference theory, ecological rationality pertains to how well people solve specific problems such as finding food, navigating through a landscape, choosing cooperative partners, or making profitable trades, and it is ultimately grounded in biological fitness (Cosmides & Tooby, 1994; Todd & Gigerenzer, 2007). Hence, a player's performance is assessed based on success in accomplishing ecologically relevant goals. Consistency plays no special role, and in fact, many high-performance decision algorithms are inconsistent in certain ways. Some inconsistency may even be inevitable in complex minds with a multiplicity of interacting psychological systems (Barrett & Kurzban, 2006; Kurzban, 2010). In vision, for instance, people's perceptions of color are inconsistent (e.g., the same patch appears blue or yellow depending on the surrounding colors), and this is due to inevitable tradeoffs when the

visual system's algorithms need to guess an object's reflectance even though it is confounded with the surrounding illumination (Purves & Lotto, 2003). In sum, our approach to performance differs from some strands of behavioral economics and is more aligned with ecological rationality, along with many other fields including perception, evolutionary biology, cognitive science, and engineering in which researchers assess performance relative to specific goals such as accuracy, fitness, health, safety, search time, object recognition, and so on.

The Orange Game

We designed an interactive online game in which a participant looks for resources to survive in a harsh environment. (A demo with instructions is available at: pdescioli.com/savingsgamedemo/game.html, see Supplemental Materials.) The more days they survive, the more real money they earn (5 cents per day). The player is stranded on an island where they consume oranges to stay alive. They start with 300 health points that decline -50 points each day due to metabolism and can be replenished by consuming oranges. If health reaches zero, then the player dies and the game is over.

Each day, a player looks for food and if they find some, they decide how much to consume and how much to save for the future. The player's income is uncertain. In the basic game, a player has a 50% chance of finding 10 oranges and a 50% chance of finding 0 oranges. When a player consumes oranges, they have diminishing returns to health each day due to constraints on digestion: the first orange adds 10 health points, the second adds 9 health points, the third adds 8 health points, and so on. These returns reset at the start of each day.

The orange game recreates a fundamental tradeoff in saving decisions. Within a given day, consuming more oranges yields diminishing returns to the player's health, so it can pay to

save oranges for future days when no oranges are found. But, if a player saves too many oranges, then they will miss opportunities to add more to their health until it is too late.

The average-plus heuristic

We use computer simulation to develop a performance benchmark for spending decisions in the orange game. (See Supplemental Materials for a simulation app and code.) We examine a set of simple heuristics for consumption. Each heuristic specifies a number of oranges k to consume when available from finding or previous savings. In addition, the heuristic checks if the initial level of consumption adds enough health to survive the current day's metabolism, and if not, it consumes additional oranges until it either can survive or runs out. We refer to this supplemental consumption as a k -plus heuristic.

We ran 10,000 simulations for each heuristic ($k = 1, 2, \dots, 10$) in two environments (see Supporting Information for simulation code). The ten-zero environment represents a high variance income where each day a player finds either 10 or 0 oranges with equal chances. The seven-three environment represents a lower variance income of 7 or 3 oranges with equal chances. Across both environments, the long-term average is held constant at 5 oranges per day.

Figure 1 shows the results. In the ten-zero environment, the 5-plus heuristic performed the best and survived $M = 25.7$ days ($SD = 13.4$). In the seven-zero environment, the 5-plus heuristic again performed the best, surviving $M = 28.5$ days ($SD = 5.6$). We refer to this high-performance decision rule as the *average-plus heuristic*, since five oranges is a player's long-term average income in the game.

Figure 1 also shows that the average-plus heuristic outperforms random consumption, in which a player randomly consumes between 0 and the number available (up to a maximum of

10, since additional oranges return 0 health). Random consumption survived for $M = 13.6$ days ($SD = 4.5$) in the ten-zero environment and $M = 15.5$ days ($SD = 3.2$) in the seven-three environment. We also checked how many oranges the average-plus heuristic actually consumed when the player had more than 5 oranges available, since supplemental consumption could vary depending on health levels. It consumed $M = 5.67$, $SD = 1.04$ oranges in the ten-zero condition and $M = 5.32$, $SD = 0.78$ oranges in the seven-three condition.

Overall, these simulations show that the average-plus heuristic is a high-performing decision rule for consumption in the game. As we will see, it also outperformed our participants in the same game which also supports its use as a benchmark. Moreover, we note that the performance of average-plus does not depend on risk preferences, since it generally performs better both at the mean and in the wider distribution of survival (see standard deviations in Figure 1). More generally, the average-plus heuristic broadly resembles prominent theories about saving from economics, which hold that rational consumers will smooth consumption over time by saving in times of plenty and consuming savings in hard times, bringing consumption closer to the long-term average income (Carroll, 2001; Modigliani & Brumberg, 1954; Friedman, 1957).

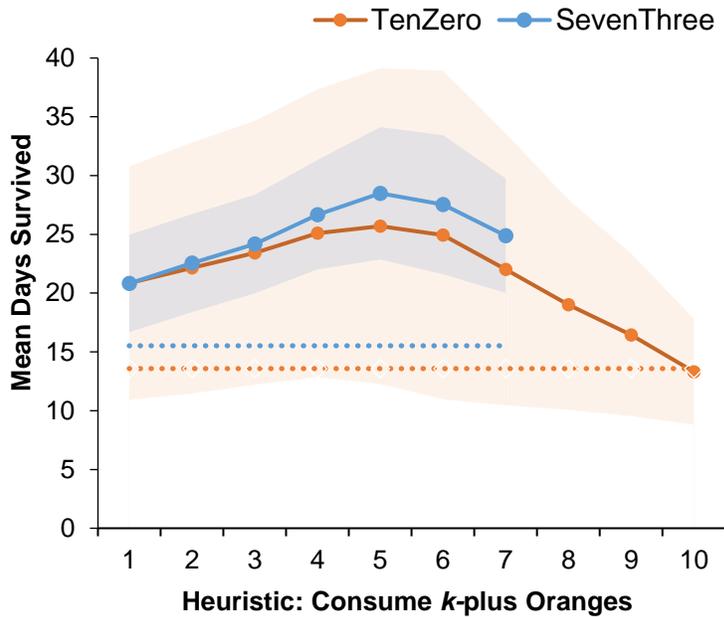


Figure 1. The mean days survived by k -plus consumption heuristics ($k = 1, 2, \dots, 10$) in the ten-zero and seven-three environments. The shaded regions are standard deviations to describe the distributions of survival. The dashed lines show survival for random consumption in each environment.

Experiment 1

In Experiment 1, we use the orange game to study participants' spending and saving decisions. We particularly focus on within-subject comparisons between how a participant spends new resources compared to how they spend stored resources. We also vary the magnitude of income variance and evaluate participants' performance relative to the average-plus heuristic.

Methods

We recruited 237 participants (43% female; age: $M = 32.3$, $SD = 12.0$) online using Amazon's Mturk (Berinsky, Huber & Lenz, 2012). We chose this sample size in advance to have

sufficient power to detect medium effect sizes in comparisons across conditions. Participants read the game instructions, which they could also consult at any time during the game. The instructions explain all aspects of the game including health, metabolism, and the diminishing returns from consuming oranges each day. Participants read that they would earn an additional 5 cents for each day they survived in the game.

Then, participants played the game (Figure 2). The player starts with 300 health points and they lose 50 points per day due to metabolism. The player's health is shown numerically and with a colored bar from red (lowest health) to green (highest health). Each day, the player looks for oranges that they can consume to replenish their health. The player finds either a lot of oranges on a windfall day or few oranges on a downturn day, with equal chances of each (50%); the exact amounts vary by condition (described below). When they find oranges, a player can drag and drop the oranges either to the dish or the basket. The oranges dropped on the dish are consumed and add points to the player's health. These health points are shown inside the dish, and they begin at 10 and diminish by 1 for each additional orange consumed in a day (10, 9, 8, ..., 0). The oranges dropped in the basket are stored for the future. At any time, a player can drag oranges from the basket to the dish to consume them (oranges do not decay).

This environment naturally provides repeated within-subject observations of a participant's spending on windfall days versus downturn days. This allows us to examine participants' spending of new resources compared to stored resources, holding constant other economically relevant factors. Hence, our primary independent variable is the within-subject manipulation of a participant's recent income, windfall or downturn.

Additionally, we manipulate between-subjects a key factor in economic theories about saving: the variance in income over the course of the game.¹ In the ten-zero condition, a player finds 10 (windfall day) or 0 oranges (downturn day) with equal chances (50% each). In the seven-three condition, a player finds 7 or 3 oranges with equal chances. We designed these values so that the average income is constant (5 oranges) while the variance is greater in the ten-zero condition. Thus, overall the experiment has a 2 (within-subject: recent income) x 2 (between-subject: income variance) mixed factor design.

A participant played the game until they died if health reached 0 or they survived the maximum of 30 days. When the game ended, participants completed a survey with demographic items and provided any general comments. Participants earned 50 cents for completing the study plus their payoff from the game, which was 5 cents per day survived. Participants' game payoffs were $M = \$1.06$, $SD = \$0.34$ in addition to the 50 cents for completing the study for a total average payment of ~\$1.50.

¹ We unintentionally recruited more participants in the ten-zero condition ($n = 137$) than the seven-three condition ($n = 98$) due to a software error.



Figure 2. The orange game.

Results

Survival. In the ten-zero condition, participants survived on average for 19.1 days, $SD = 7.2$, and 16% of participants lasted the maximum 30 days. To compare participants' survival to the simulated average-plus heuristic, we first capped the simulated survival amounts to a maximum of 30 days to match the experiment. Participants survived significantly less days than the average-plus heuristic ($M = 22.0$ days, $SD = 7.8$), $t(10,135) = 5.78$, $p < .001$, and they survived more days than random consumption ($M = 13.6$ days, $SD = 4.4$), $t(10,135) = 11.75$, $p < .001$.

In the seven-three condition, participants survived on average for 23.3 days, $SD = 5.1$, and 22% lasted the maximum 30 days. Participants survived significantly less days than the average-plus heuristic ($M = 26.9$ days, $SD = 3.7$), $t(10,096) = 9.70$, $p < .001$, and significantly more days than random consumption ($M = 15.5$ days, $SD = 3.2$), $t(10,096) = 23.56$, $p < .001$.

Comparing across conditions, participants survived more days in the seven-three condition than the ten-zero condition, $t(233) = 6.03, p < .001$.

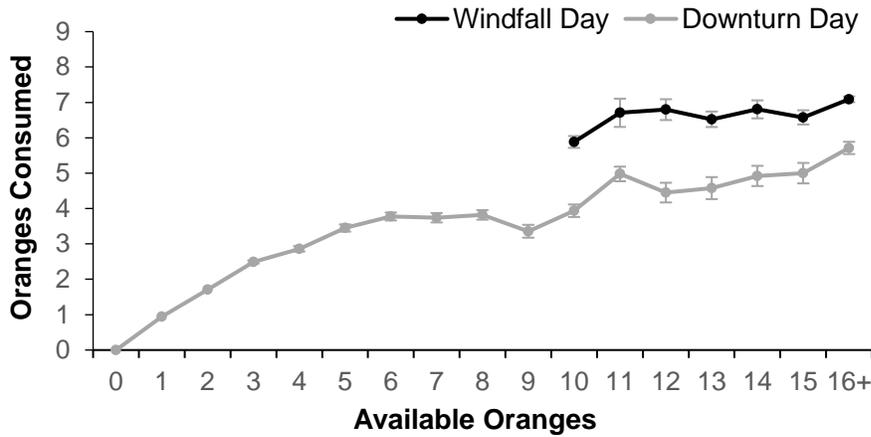
Consumption of new resources versus stored resources. To test the resource preservation hypothesis, we examine how participants consumed new oranges compared to stored oranges. Importantly, the distinction between new and stored oranges is irrelevant for survival, the only quantity that can affect performance is a player's *available oranges*, which is new and stored oranges combined. In economic terms, oranges are fungible and the labels of new versus stored are irrelevant for optimal usage, analogous to mental accounts for current income versus current wealth (Thaler, 1999). Hence, we can test for the effects of preservation motives by comparing consumption when a player has the same available oranges but in different combinations of new and stored oranges.

For instance, from a performance perspective, a player who found 0 new oranges and has 10 in savings is in an equivalent position to a player who oppositely found 10 new oranges and has 0 in savings, because both players have 10 available oranges. But from a preservation perspective, a player will be more willing to spend recently found oranges than to deplete their reserves, so they will spend less on downturn days when they receive less new oranges.

Figure 3 shows participants' consumption by their available oranges (new plus stored oranges), and also by whether it was a windfall day or downturn day, meaning they found 10 or 0 new oranges, respectively. We can compare vertically at each level of available oranges to see how participants' consumption depended on their recent income. The figure shows that in both conditions participants consistently consumed less oranges on downturn days than windfall days, holding constant the available oranges. In other words, the sizeable gap between the windfall and

downturn values violates the economic principle of fungibility, which requires that spending should be the same when the available oranges are the same. Instead, participants were more willing to consume the new oranges from a windfall and less willing to consume stored oranges after a downturn.

A. Ten-zero condition



B. Seven-three condition

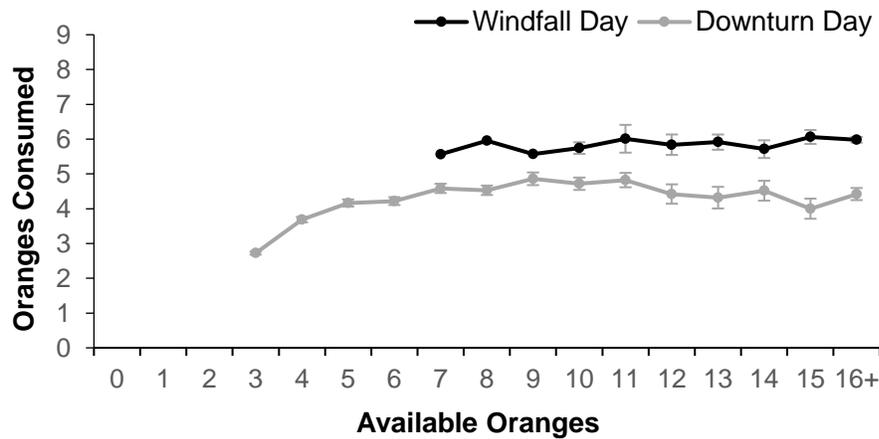


Figure 3. The mean (*SE*) number of oranges a participant consumed by the number of oranges that were available (new plus stored oranges). At each level of available oranges, the participant's supply of oranges differs only in the relative proportion of new oranges (greater on windfall days) and the proportion of stored oranges (greater on downturn days), which is an irrelevant difference for efficient consumption because new and stored oranges are fungible. Consumption is shown for windfall and downturn days in the ten-zero condition (panel A) and the seven-three condition (panel B).

To analyze these effects, we use regression to estimate how consumption depends on the recent income (windfall or downturn) and income variance throughout the game (ten-zero or seven-three). To capture within-subject effects, we include a random effect for participant in the model. To control for available oranges, we focus the analysis on days when participants had at least five oranges available; this holds constant that a player had sufficient oranges to consume the long-term average of five, independent of fluctuations in their recent income.

Table 1 shows the results. We find a significant effect of downturns, showing that in the seven-three condition participants consumed 1.15 less oranges on downturn days than windfall days, Wald $\chi^2(1) = 169.92, p < .001$. Moreover, the effect of downturn plus its interaction with ten-zero shows that participants also consumed significantly less oranges after a downturn (1.89 less; test of combined coefficients for downturn and ten-zero*downturn condition) in the ten-zero condition, Wald $\chi^2(1) = 500.44, p < .001$. Finally, the significant interaction term shows that the effect of downturns was greater when downturns were more severe in the ten-zero condition than the seven-three condition.

These results show that a performance-irrelevant factor—whether a participant’s resources were newly found on a windfall day or located in storage on a downturn day—systematically affected consumption, and more so when income had greater variance. Consistent with the preservation hypothesis, participants consumed less on a downturn day when they had to deplete their stored resources in order to smooth consumption.

Table 1
Regression of Consumption by Recent Income and
Income Variance

	Oranges Eaten per Day
Downturn	-1.15 (0.09)
Ten-zero	0.72 (0.15)
Ten-zero * Downturn	-0.74 (0.12)
Constant	5.70 (0.11)

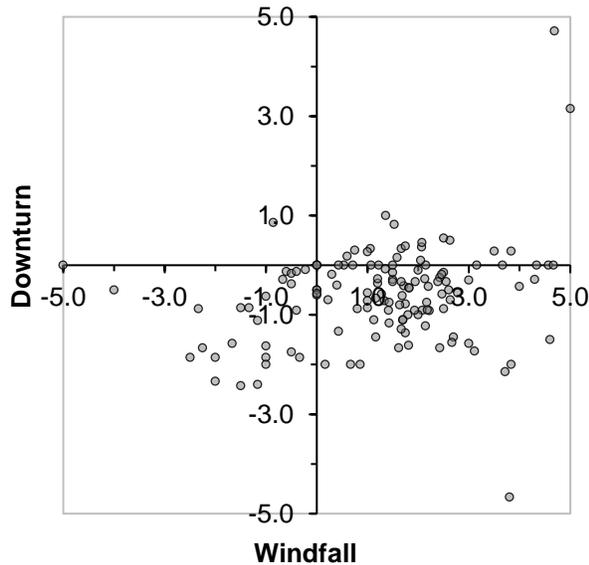
Note. The reference category is a windfall day in the seven-three condition. The model includes a random effect for participant. The analysis includes consumption decisions when a participant has more than five oranges available. Standard errors are shown in parentheses. All coefficients are statistically significant, $ps < .001$.

Performance compared to the average-plus heuristic. We next examine participants' performance compared to the average-plus heuristic. This heuristic survived the most days in simulations for both the seven-three and ten-zero environments (see Introduction), and it also survived longer than participants, supporting its use as a performance benchmark.

We measure each participant's efficiency with separate scores for windfall and downturn days. Each efficiency score is based on the difference between a participant's consumption and what the average-plus heuristic would consume in the same situation. Hence, a score of 0 represents optimal efficiency while negative and positive values represent under- and overconsumption, respectively. We average these differences across all of a participant's decisions, separately for windfall and downturn days. This measure captures a participant's ability to consume efficiently to survive in the game.

Figure 4 shows the distribution of participants' efficiency scores in each condition. Comparing across conditions, participants' consumption appears closer to the average-plus heuristic (0,0) in the seven-three condition than the ten-zero condition. In both conditions, the modal pattern was consuming too much on windfall days and too little on downturn days (52% of participants in the ten-zero condition and 68% in the seven-three condition), whereas the reverse pattern rarely occurred (1% in ten-zero condition and 0% in seven-three condition). Moreover, the greater spread in scores in the ten-zero compared to seven-three condition suggests that greater variance in income led to greater variability in efficiency across individuals.

A. Ten-zero condition



B. Seven-three condition

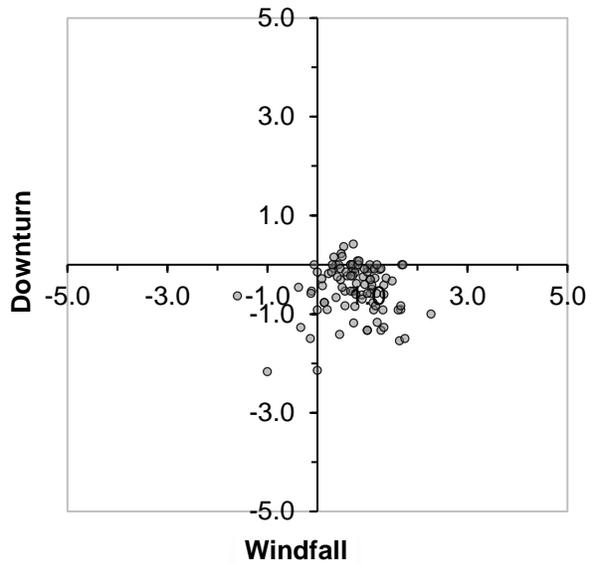


Figure 4. Participants' efficiency scores on windfall and downturn days. The scores are the average difference from the average-plus heuristic on windfall and downturn days. Scores are shown for the ten-zero condition (panel A) and the seven-three condition (panel B). The axes are truncated at 5 and -5 for presentation (2% of participants were outside of this range).

We next examine more closely how recent income and income variance affected participants' performance. Figure 5 shows the mean efficiency scores on windfall and downturn days by income variance. Irrespective of income variance, we find that participants consumed significantly more than the average-plus heuristic (represented by 0) on windfall days and significantly less on downturn days. Moreover, we find that overspending on windfall days was greater when resources were more variable in the ten-zero condition than in the seven-zero condition, $t(233) = 2.21, p < .05$, whereas participants' underspending after downturns remained consistent and did not differ between ten-zero and seven-three conditions, $t(233) = 0.27, p = .79$.

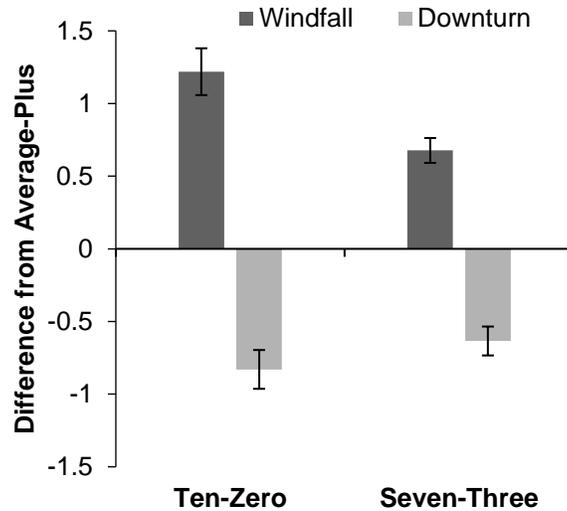


Figure 5. Mean (SE) efficiency scores by windfall day and downturn day. A participant's efficiency score is the average difference between the number of oranges they consumed and the number that the average-plus heuristic would consume. All means significantly differ from zero ($p < .01$), where zero represents matching the high-performance average-plus heuristic.

To further analyze efficiency scores, we conducted a regression with predictors for recent income, income variance, the interaction, and a random effect for participant. Table 2 shows the results. The significant interaction shows that the effect of downturns differed across the ten-zero and seven-three conditions. In the seven-three condition, participants consumed more than the heuristic on windfall days (0.66; Wald $\chi^2(1) = 17.51, p < .001$) and less on downturn days (-0.69; test of combined coefficients for constant and downturn; Wald $\chi^2(1) = 19.38, p < .001$). In the ten-zero condition, participants consumed more than the heuristic on windfall days (1.16, test of combined coefficients for constant and ten-zero condition; Wald $\chi^2(1) = 76.18, p < .001$) and less than the heuristic on downturn days, (-0.64; test of all combined coefficients; Wald $\chi^2(1) = 23.71, p < .001$).

Table 2

Regression of Efficiency Scores by Recent Income and
Income Variance

	Difference from Average-Plus
Downturn	-1.35 (0.17)
Ten-zero	0.50 (0.21)
Ten-zero * Downturn	-0.45 (0.22)
Constant	0.66 (0.16)

Note. The reference category is efficiency on windfall days in the seven-three condition. The model includes a random effect for participant. Standard errors are shown in parentheses. All coefficients are statistically significant ($p < .05$).

Discussion

In sum, we find that participants consumed more on windfall days with an abundance of new oranges, and they consumed less on downturn days when consumption required depleting their stored oranges; this difference occurred while holding constant the available oranges (new plus stored) which is the only economically relevant quantity since new and stored oranges are fungible and equivalent. Second, when we measured performance compared to the average-plus heuristic, we found that participants not only spent too much after windfalls, they also spent too little after downturns. Hence, supporting the resource preservation hypothesis, participants

overreacted to their recent income in both directions. Finally, these effects of recent income were amplified when income was higher variance in the ten-zero condition compared to the seven-three condition. Consequently, participants' decisions strayed further from efficiency when their income was higher variance.

Experiment 2

In Experiment 2, we look closer at how participants spend in hard times. In Experiment 1, participants spent too little after downturns which is consistent with the idea that they are reluctant to deplete their savings. However, an alternative possibility is that underspending was a byproduct of windfall effects. For instance, participants who received a windfall in income might have perceived a subsequent downturn as more severe by contrast. Hence, it could be that it was mainly the contrast with windfalls that caused decreases in spending rather than the downturn itself.

To further examine underspending, the present experiment tests whether it is possible to dissect and separate underspending from overspending. If people's motive to preserve resources shapes underspending, then it should be possible to separately amplify overspending or underspending by varying the frequency of windfalls or downturns. In the previous experiment, participants' income fluctuated up or down each day, randomly alternating windfalls or downturns. In the current experiment, participants instead receive a steady income most periods (75% chance) and then they only occasionally (25% chance) receive either windfalls or downturns, depending on whether they are assigned to the downturn condition or windfall condition (varied between-subjects). This allows us to control and manipulate whether participants experience occasional windfalls or occasional downturns.

The preservation hypothesis predicts that participants will still underspend after downturns, even when they occur in isolation from windfalls of equivalent size. Similarly, preservation and mental accounting predict overspending after windfalls in the windfall condition, even when windfalls are isolated from downturns. The preservation hypothesis predicts overspending in this case because a windfall is a cue for a period of abundance and preservation motives are less intense for recent income that has not yet been stored. Moreover, these predictions imply that the greatest underspending will occur after downturns in the downturn condition and the greatest overspending will occur after windfalls in the windfall condition.

To test these ideas and evaluate participants' performance, we again compare their choices to the average-plus heuristic. We checked and found that this heuristic continued to perform best in simulations under the new parameters of the windfall and downturn conditions (see Supplemental Materials).

Methods

We recruited 158 participants on Mturk (47% female; age: $M = 34.8$, $SD = 10.7$). Participants played a variant of the orange game (a demo is available at: pdescioli.com/orangeGame/orange.dt.demo.html) with the same general procedures as in Experiment 1. Participants' game payoffs were $M = \$1.09$, $SD = \$0.33$ in addition to the 50 cents for completing the study for a total average payment of $\sim \$1.50$.

Participants were assigned to the windfall condition ($n = 80$) or downturn condition ($n = 78$). In the windfall condition, each day a player finds 3, 3, 4, or 10 oranges with equal chances (25% chance each), and in the downturn condition, a player finds 7, 7, 6, or 0 oranges with equal

chances. We designed these values so that the average income is constant (5 oranges) across conditions while the distribution of income varies. Specifically, a player receives a steady stream of income near the long-term average most of the time (75% of days), and then this typical income is occasionally (25%) punctuated with either windfalls (windfall condition) or downturns (downturn condition).

Each game provides repeated within-subject observations of a participant's consumption on typical days with near-average income compared to days with a transitory income shock. And between-subjects, we manipulate whether the shocks are windfalls or downturns. Thus, overall the experiment has a 2 (within-subject: typical income or income shock) x 2 (between-subject: shocks are windfalls or downturns) mixed factor design.

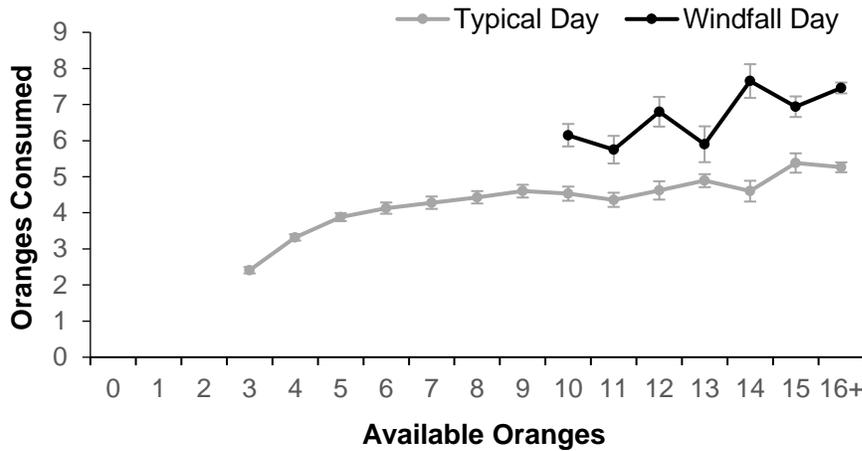
Results

Survival. In the windfall condition, participants survived on average for 22.3 days, $SD = 6.0$, and 23% of participants lasted the maximum 30 days. Participants survived less days than the average-plus heuristic ($M = 25.6$ days, $SD = 4.7$), $t(10,078) = 6.46$, $p < .001$, and they survived more days than random consumption ($M = 15.1$ days, $SD = 3.4$), $t(10,078) = 18.71$, $p < .001$. In the downturn condition, participants survived on average for 23.0 days, $SD = 6.4$, and 22% lasted the maximum 30 days. Participants survived less than the average-plus heuristic ($M = 25.2$ days, $SD = 5.6$), $t(10,076) = 3.47$, $p < .001$. In addition, participants survived more days than random consumption ($M = 15.0$ days, $SD = 3.7$), $t(10,076) = 18.73$, $p < .001$. Comparing across conditions, participants' survival did not significantly differ between the downturn and windfall conditions, $t(156) = 0.72$, $p = .47$.

Consumption of new resources versus stored resources. As in Experiment 1, we first examine the effect of ownership by comparing at each level of available oranges (new plus stored) how many oranges participants consume when they recently received a typical income or an income shock (windfall or downturn). Recall that efficient spending depends on the total available oranges rather than the recent income of new ones. So, an efficient player will spend no differently on days with typical income versus an income shock, after holding constant the total available oranges. In contrast, if a player attributes greater value to stored resources, then they will spend more after a windfall than with typical income, and less after a downturn than with typical income. This is because participants can consume more after a windfall without depleting their savings, whereas consuming the same amount after a downturn requires depleting their previous reserves.

Figure 6 shows participants' mean consumption per day by the number of available oranges (new plus stored), and also by whether it was a typical day versus a day with a windfall or downturn. In both conditions, participants' consumption depends on how much they found that day rather than only on the total amount they had available, again violating the principle of fungibility. In the windfall condition, participants consumed more oranges on windfall days than on typical days. In the downturn condition, participants consumed less oranges on downturn days than on typical days.

A. Windfall condition



B. Downturn condition

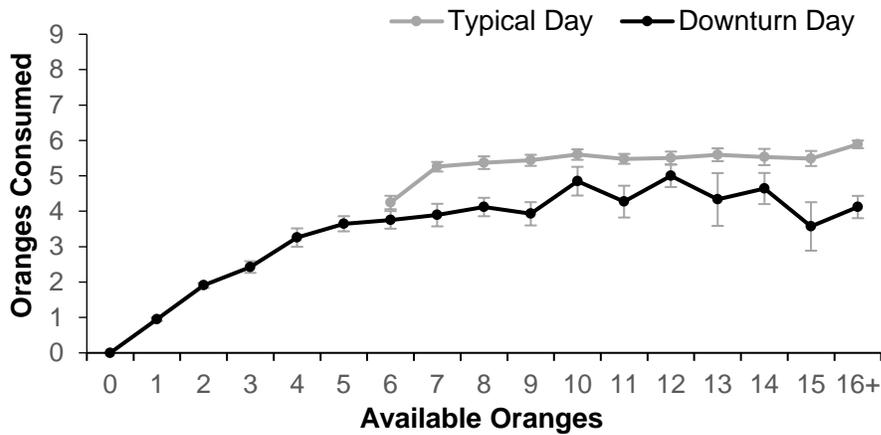


Figure 6. The mean (*SE*) number of oranges a participant consumed by the number of oranges that were available (new plus stored oranges). Consumption is shown for the windfall condition (panel A) on typical days (found 3 or 4, 75% of days) and windfall days (found 10, 25% of days), and the downturn condition (panel B) on typical days (found 6 or 7, 75% of days) and downturn days (found 0, 25% of days).

To look further, we use regression to estimate how consumption depends on an income shock and the direction of the shock (windfall or downturn). We focus the analysis on days when

participants had more than five oranges available, thus holding constant that a player had sufficient oranges to consume the long-term average of five. Table 3 shows the results. We find a significant effect of an income shock, showing that in the windfall condition, participants consumed 2.17 more oranges on windfall days, Wald $\chi^2(1) = 447.97, p < .001$. In the downturn condition, participants consumed 1.15 less oranges after a downturn, Wald $\chi^2(1) = 86.71, p < .001$ (test of combined coefficients for shock and shock*downturn condition). These results show that a performance-irrelevant factor—whether or not a participant had an income shock on a given day—systematically affected consumption according to the direction of the shock, namely spending more after a windfall and less after a downturn.

Table 3

Regression of Consumption by Income Shock and Shock Direction

	Oranges Consumed
Shock	2.17 (0.10)
Downturn Condition	0.74 (0.14)
Shock * Downturn Condition	-3.32 (0.16)
Constant	4.54 (0.10)

Note. The reference category is a typical day in the windfall condition. Model includes random effect for participant. The analysis includes consumption decisions when a participant has at least five available

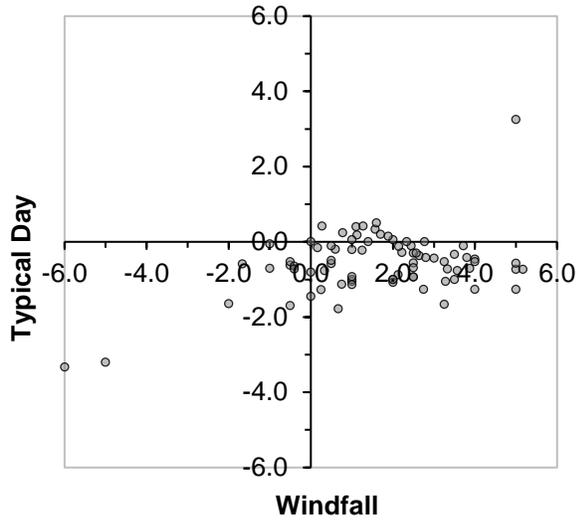
oranges (the long-term average). Standard errors are shown in parentheses. All coefficients are statistically significant, $ps < .001$.

Performance compared to the average-plus heuristic. We next examine participants' performance compared to the average-plus heuristic (see Supplemental Materials for simulations). We measure each participant's efficiency with separate scores for typical days and days with shocks (windfall or downturn days). As in Experiment 1, the efficiency score is based on the difference between a participant's consumption and what the heuristic would consume in the same situation, hence 0 represents optimal efficiency while negative and positive values represent under- and overconsumption, respectively.

Figure 7 shows the distribution of efficiency scores in each condition.² In the windfall condition, the modal pattern (64% of participants) is consuming too much on windfall days and too little on typical days; in the downturn condition, the modal pattern (44% of participants) is consuming too little on downturn days and too much on typical days; whereas the reverse patterns rarely occurred (0% in windfall condition, 1% in downturn condition).

² In the windfall condition, 3 participants died before receiving a windfall so they are not included in the efficiency analysis.

A. Efficiency scores, windfall condition



B. Efficiency scores, downturn condition

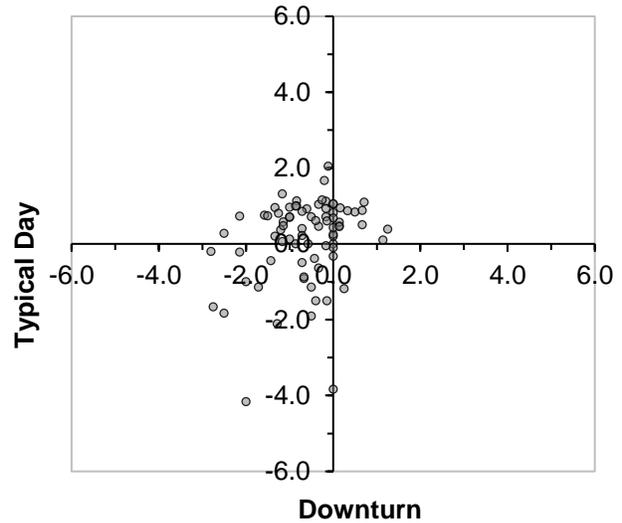


Figure 7. Participants' efficiency scores on typical, windfall, and downturn days. The scores are the average difference from the average-plus heuristic on typical, windfall, and downturn days. Scores are shown for the windfall condition (panel A) and the downturn condition (panel B).

We next examine more closely how an income shock and its direction affected participants' performance. Figure 8 shows the mean efficiency scores on days with shocks and typical days, separately for the windfall and downturn conditions. In the windfall condition, participants consumed significantly more than the average-plus heuristic (represented by 0) on windfall days and significantly less than optimal on typical days. In the downturn condition, participants consumed significantly less than the average-plus heuristic on downturn days and did not differ from the average-plus heuristic on typical days.

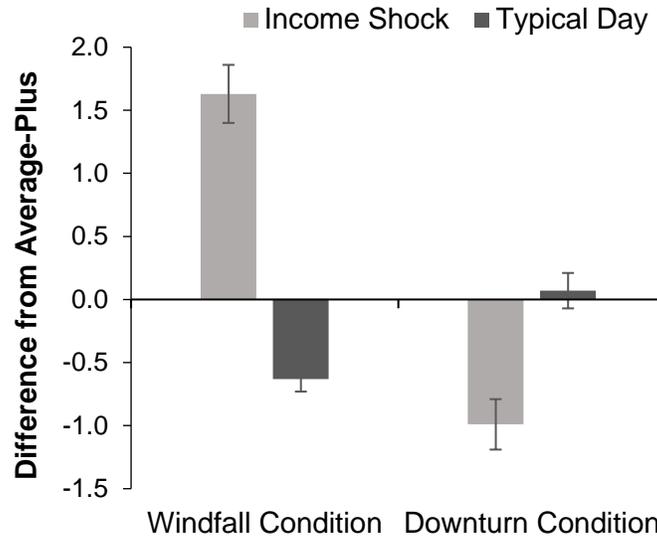


Figure 8. Mean (SE) efficiency scores by their income level that day, which was typical or windfall in the windfall condition and typical or downturn in the downturn condition. A participant's efficiency is the average difference between their consumption and what the average-plus heuristic would consume. All means significantly differ from zero ($p < .01$), except for the typical day in the downturn condition ($p = .64$).

We conducted a regression of efficiency scores with predictors for income shock, shock direction, the interaction, and a random effect for participant. Table 4 shows the results. The significant interaction shows that the effect of a shock differed across the windfall and downturn conditions. In the windfall condition, participants consumed more than the heuristic on windfall days (1.61; Wald $\chi^2(1) = 106.48, p < .001$; test of combined coefficients for constant and shock) and less than the heuristic on typical days (-0.63; Wald $\chi^2(1) = 16.49, p < .001$).

In the downturn condition, participants' consumption did not differ from the average-plus heuristic on typical days (0.07; Wald $\chi^2(1) = 0.18, p = .67$; test of combined coefficients for

constant and downturn condition) and they consumed significantly less than optimal after downturns (-0.73; Wald $\chi^2(1) = 21.13, p < .001$; test of all coefficients combined).

Table 4

Regression of efficiency scores by income shock and direction of shock

	Difference from Average-Plus
Shock	2.24 (0.18)
Downturn Condition	0.69 (0.22)
Shock * Downturn Condition	-3.03 (0.25)
Constant	-0.63 (0.15)

Note. The reference category is a typical day in the windfall condition. Model includes random effect for participant.

Standard errors are shown in parentheses. The constant term is not significant, and the remaining coefficients are significant, $ps < .001$.

Discussion

We find, first, that participants spent more on windfall days than typical days, and they spent less on downturn days than typical days; these comparisons hold constant the available oranges so they represent violations of fungibility. Second, when we measured participants' performance compared to the average-plus heuristic, we found that participants spent too much after windfalls and they spent too little after downturns. These findings show that participants overreacted to their recent fortunes in both directions.

Moreover, these results show, as predicted by the resource preservation hypothesis, that underspending after downturns can be separated and isolated from overspending after windfalls. It reveals how the particular pattern of an income stream—in this case steady income punctuated by either windfalls or downturns—shapes whether participants overspend or underspend.

Experiment 3

In Experiment 3, we change the theme of the game to managing money in an urban household. We otherwise use the same underlying payoffs and experimental design as in Experiment 2. This allows us to test for windfall and downturn effects when the game is framed in terms of managing money in a household rather than foraging for survival.

Methods

We recruited 159 participants (50% female; age: $M = 32.6$, $SD = 12.5$) on MTurk. Participants played the big city game (Figure 9), which is the same as the game in Experiment 2 except for the framing. The player is a worker who recently moved to the big city and needs to earn a living to make it in the city. The player starts with 300 health points and they lose 50 points per month due to their accumulating needs for food, housing, and other necessities. Each month, the player looks for work to earn money that they can spend at the shop to replenish their health. When they find a job and earn money, the player can drag and drop dollars either to the shop or the piggy bank. Dollars dropped on the shop are spent to add points to the player's health. Dollars dropped in the piggy bank are stored for the future.

Participants were assigned to the downturn condition ($n = 79$) or windfall condition ($n = 80$). As in Experiment 2, in the windfall condition, the player earns 3, 3, 4, or 10 dollars with

equal chances (25% each); in the downturn condition, the player earns 7, 7, 6, or 0 dollars. A participant played the game until their health reached 0, which meant they had to leave the city, or when they reached the maximum of 30 months. Participants' game payoffs were $M = \$1.03$, $SD = \$0.33$ in addition to the 50 cents for completing the study for a total of $\sim \$1.50$.

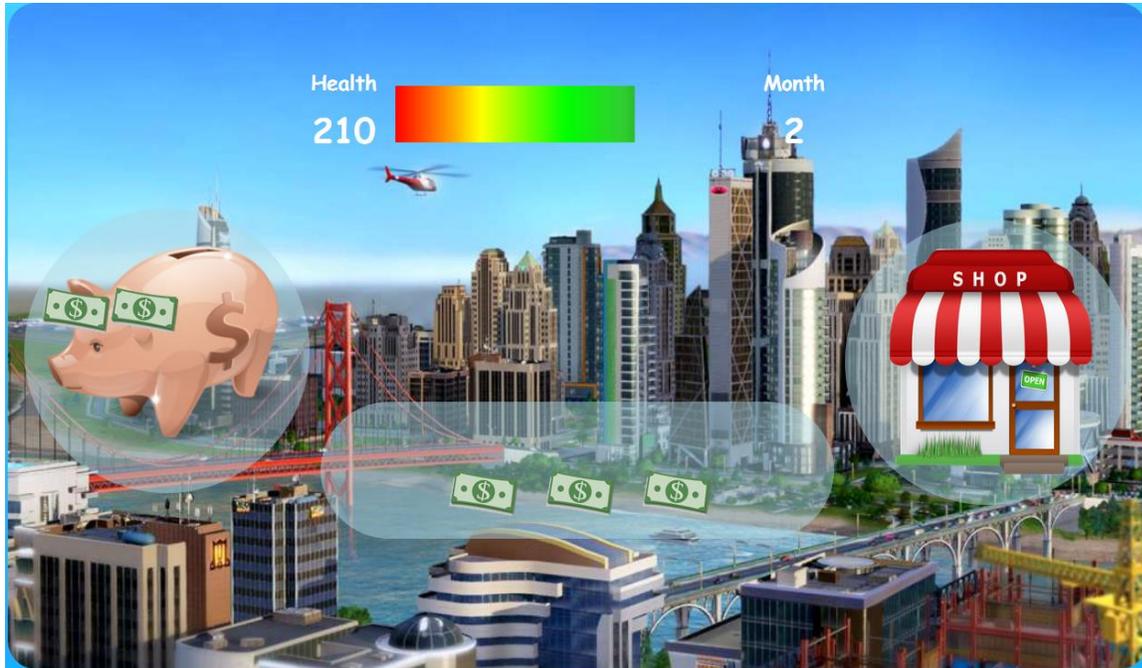


Figure 9. The big city game.

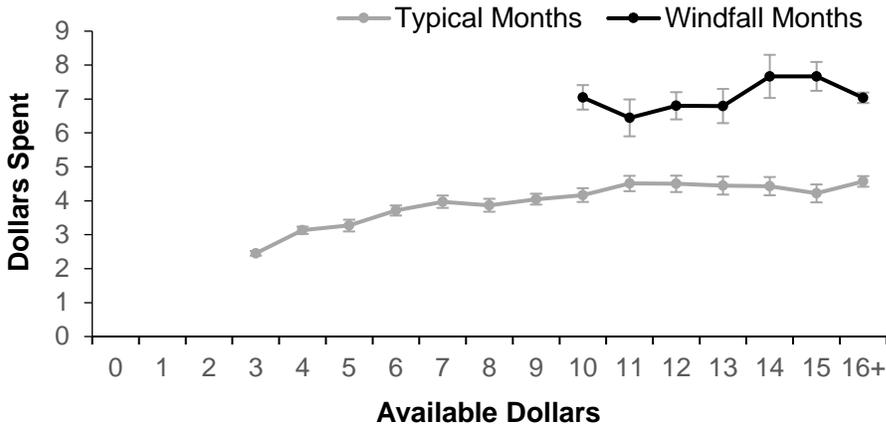
Results

We conducted the same analyses as in Experiment 2 and found the same basic patterns of results. For survival, participants' performance was similar to Experiment 2 (Supplemental Materials). For the consumption of new versus stored resources, Figure 10 shows that in the windfall condition, participants consumed more money on windfall months than on typical months (panel A), holding constant the money available; and in the downturn condition, participants spent less money on downturn months than on typical months, again holding

constant the money available (panel B). Moreover, a regression analysis confirmed that, as in Experiment 2, participants consumed more money after a windfall compared to typical months, and less money after a downturn compared to typical months, holding constant economically relevant factors (Supplemental Materials).

When we compare participants' performance to the average-plus heuristic, we found similar distributions as in Experiment 2 (Supplemental Materials). Figure 11 shows the mean efficiency scores on months with shocks and typical months, separately for the windfall and downturn conditions. In the windfall condition, participants consumed significantly more than the average-plus heuristic (represented by 0) on windfall months and significantly less than it on typical months. In the downturn condition, participants consumed significantly less than the average-plus heuristic on downturn months and did not differ from it on typical months. A regression analysis further confirmed that participants tended to overspend after windfalls and underspend after downturns (Supplemental Materials).

A. Windfall condition



B. Downturn condition

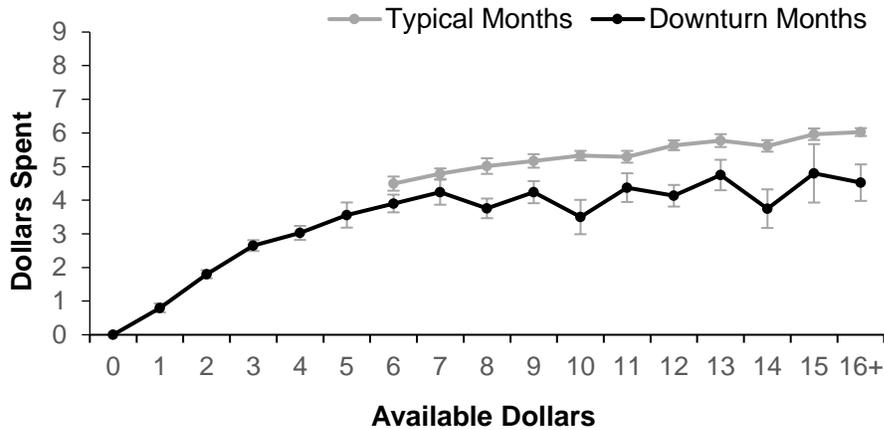


Figure 10. The mean number of dollars a participant consumed by the number of dollars that were available. Consumption is shown for the windfall condition (panel A) on typical months (found 3 or 4, 75% of months) and windfall months (found 10, 25% of months), and the downturn condition (panel B) on typical months (found 6 or 7, 75% of months) and downturn months (found 0, 25% of months).

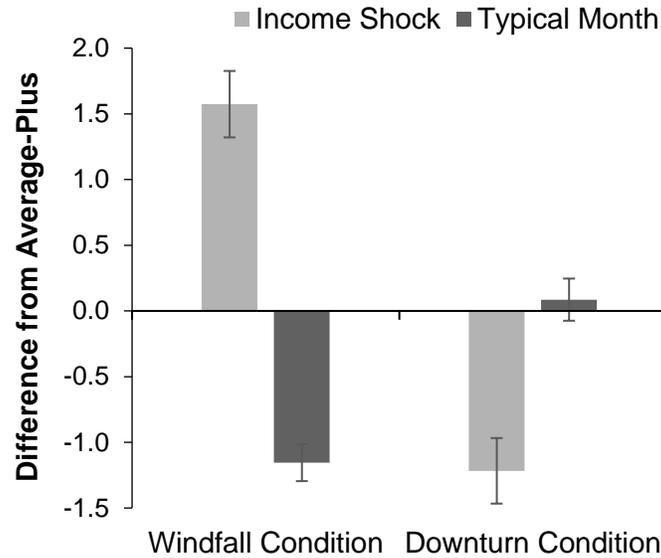


Figure 11. Mean (SE) efficiency scores by their income level that month, which was typical or windfall in the windfall condition and typical or downturn in the downturn condition. A participant's efficiency is the average difference between their consumption and the average-plus heuristic. All means significantly differ from zero ($p < .01$), except for the typical month in the downturn condition ($p = 0.58$).

Discussion

To summarize, in a game with a new theme of working for money to make it in the big city, we found the same basic patterns as in Experiment 2. First, with the available money held constant, participants spent more money on windfall months than typical months, and they spent less money on downturn months than typical months. Second, when we assessed performance compared to the average-plus heuristic, we again found that participants consistently spent too much after windfalls and too little after downturns.

Experiment 4

In Experiment 4, we change the theme of the game to managing the national budget. Otherwise, the underlying payoffs and experimental design are the same as in Experiments 2 and 3. This enables us to test for windfall and downturn effects when the game is framed in terms of managing a nation's finances.

Methods

We recruited 160 participants (50% female; age: $M = 35.04$, $SD = 13.60$) on MTurk. Participants played the national budget game (Figure 12), which is the same as the game in Experiment 2 and 3 except for the framing. The player is the president of the United States who controls the national budget with the support of the citizens and the legislature. The player starts with the nation's wellbeing at 300 points and the nation loses 50 points per month due to the citizens' accumulating needs for food, housing, and other necessities. Each month, the player collects tax revenue that they can spend to provide public services to the people to replenish their wellbeing. When they receive tax revenue, the player can drag and drop dollars either onto the people or the national treasury. Dollars dropped onto the people are spent on public services to add points to the nation's wellbeing. Dollars dropped in the national treasury are stored for the future.

Participants were assigned to the downturn condition ($n = 80$) or windfall condition ($n = 80$). The two conditions have the same payoffs as Experiments 2 and 3: In the windfall condition, the player collects 3, 3, 4, or 10 tax dollars with equal chances (25% each); in the downturn condition, the player collects 7, 7, 6, or 0 tax dollars. If wellbeing reached 0, then the player lost the support of the citizens and legislature, so they lost control of the budget and the game was over. A participant played the game until the nation's wellbeing reached 0, which meant they had

to leave office, or they kept the nation's wellbeing above 0 for the maximum of 30 months.

Participants' game payoffs were $M = \$1.06$, $SD = \$0.32$ in addition to the 50 cents for completing the study for a total of $\sim \$1.50$.

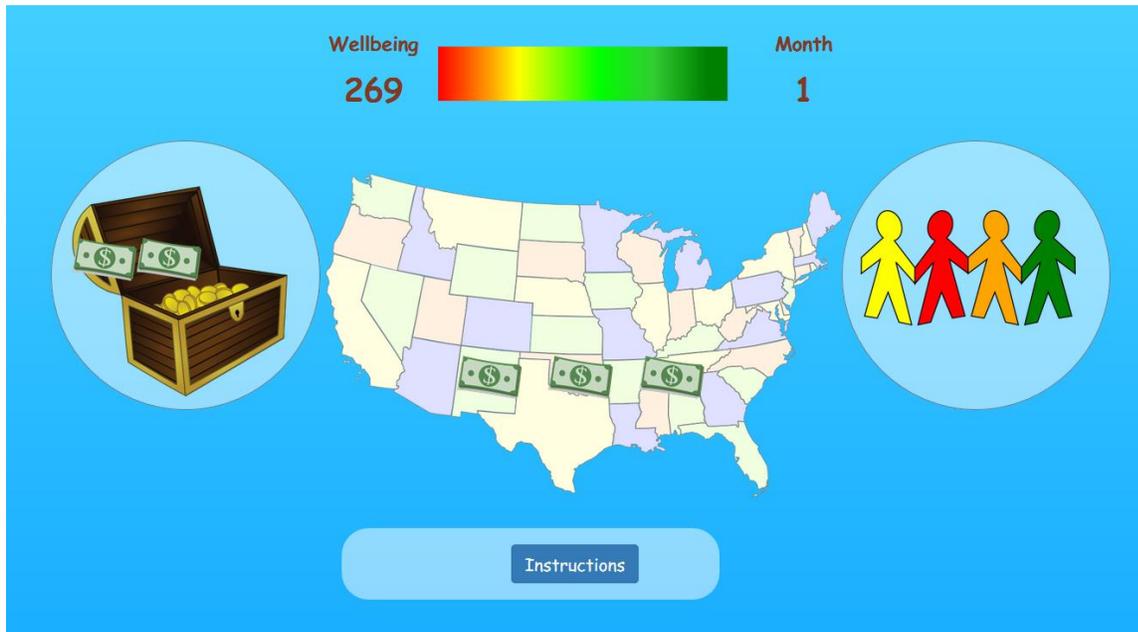


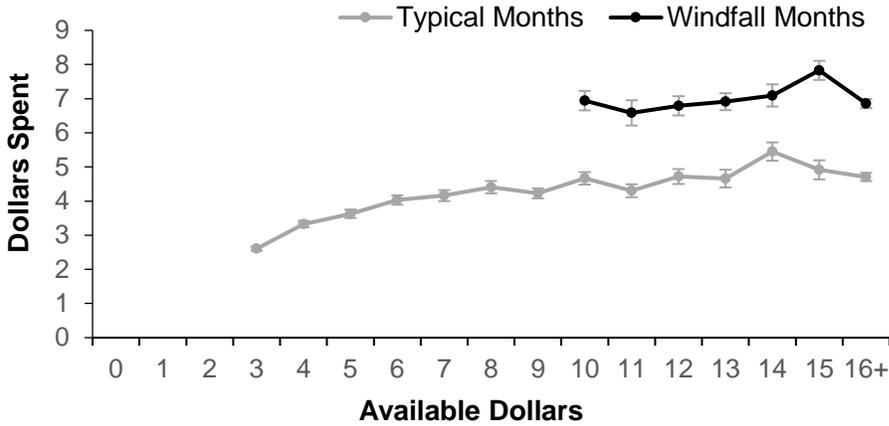
Figure 12. The national budget game.

Results

Overall, the results showed the same patterns as for the foraging and household finance themes in Experiments 2 and 3. For survival, participants' performance was similar to before (Supplemental Materials). For the consumption of new versus stored resources, Figure 13 shows that in the windfall condition, participants spent more money during windfall months than typical months (panel A), holding constant the available money; in the downturn condition, participants spent less money during downturn months than typical months, holding constant the available money (panel B). A regression analysis confirmed that, as in Experiment 2 and 3, participants

spent more money after a windfall compared to typical days, and less money after a downturn, holding constant economically relevant factors (Supplemental Materials).

A. Windfall condition



B. Downturn condition

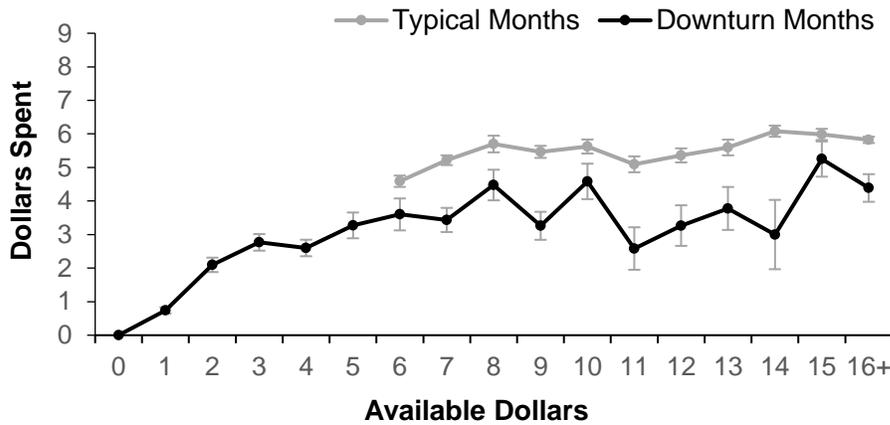


Figure 13. The mean number of dollars a participant spent by the number of dollars that were available. Spending is shown for the windfall condition (panel A) on typical months (found 3 or 4, 75% of months) and windfall months (found 10, 25% of months), and the downturn condition (panel B) on typical months (found 6 or 7, 75% of months) and downturn months (found 0, 25% of months).

Comparing participants' performance to the average-plus heuristic, we found similar distributions as in Experiment 2 and 3 (see Supplemental Materials). Figure 14 shows the mean efficiency scores on months with shocks and typical months, separately for the windfall and downturn conditions. In the windfall condition, participants spent significantly more than the average-plus heuristic (represented by 0) on windfall months and less than optimal on typical months. In the downturn condition, participants spent less than the average-plus heuristic on downturn months and did not significantly differ from zero (matching the heuristic) on typical months. A regression analysis showed again that participants tended to spend too much after windfalls and spend too little after downturns (Supplemental Materials).

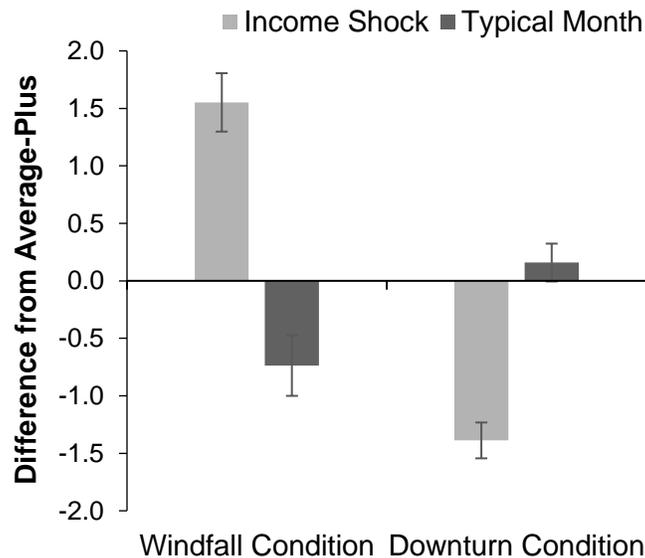


Figure 14. Mean (SE) efficiency scores by income level, which was typical or windfall in the windfall condition and typical or downturn in the downturn condition. A participant's efficiency is the average difference between their consumption and the average-plus heuristic. All means significantly differ from zero ($p < .01$), except for the typical month in the downturn condition ($p = .26$).

Discussion

In short, participants' saving decisions in a game with a political theme, managing a national budget, show the same patterns as Experiments 2 and 3. First, with the available tax dollars held constant, participants spent more money on windfall months than typical months, and they spent less money on downturn months than typical months. Second, compared to the average-plus heuristic, participants spent too much after windfalls and too little after downturns.

General Discussion

Overall, these experiments indicate that people's spending and saving decisions are not only a matter of time but also resource management. In an interactive economic game in which participants could store resources and their performance could be measured, we found that in addition to spending too much after windfalls, participants also spent too little after downturns. This observation points to the existence of a psychological motive that can counteract temporal discounting, and particularly a motive that is more active during a downturn in income. The resource preservation hypothesis proposes that the missing motive is supplied by psychological mechanisms for managing resources, which guide people to accumulate, manage, preserve, and defend reserves of valuable resources (Boyer, 2015; DeScioli & Wilson, 2011; Stake, 2004). During a downturn, a person's motive to consume is opposed by a motive to preserve their savings rather than see their wealth vanish in consumption. In many cases, this preservation motive could help to stretch supplies through a long period of scarcity, but it could also lead to underspending when scarcity is actually transitory, like in the present experiments and some real-world cases of unemployment, natural disasters, and other hardships.

Further, the results of Experiment 2-4 show that the downturn effect can be separated from the windfall effect by manipulating a participant's stream of income. This confirms an additional prediction of the preservation hypothesis, which is that people decrease spending after downturns per se rather than only when downturns follow large windfalls.

More generally, participants' consumption decisions systematically violated the economic principle of fungibility: Rather than treating new and stored resources the same, participants consumed more new resources and less stored resources, holding constant their available resources (new plus stored). This observation is consistent with the theory of mental accounting (Thaler, 1985, 1999), which maintains that people's spending depends on how they label and categorize wealth into different mental accounts. Furthermore, the present experiments suggest that people can readily form different mental accounts even within a stylized economic game. Participants could have viewed the whole game as one task with earnings that fall into a single mental account, but instead they treated new income differently from stored wealth.

The finding that participants spent too little after downturns is consistent with the large economic literature on consumer spending, which has similarly found reduced spending after transitory downward shocks to income (reviewed in Jappelli & Pistaferri, 2010). The psychology of resource preservation could help understand these and other spending decisions. For instance, another puzzling observation is that many consumers hold money in low-yield savings while simultaneously holding high-interest debt, rather than using their savings to pay the debt (Gross & Souleles, 2002). This costly inefficiency can be understood in terms of people's different mental accounts, and particularly, a motive to preserve savings even when it could be more efficient to pay off the debt. In one study, some participants said that they would pay for an emergency with a high-interest credit card instead of their savings, and they were more likely to

do so when the savings had been set aside for a responsible goal (supporting their children) compared to an ordinary goal (buying a car), which illustrates how mental accounts can lead to an inefficient use of savings (Sussman & O'Brien, 2016).

These patterns of consumer spending have been a longstanding source of debate with many alternative interpretations. We suggest that economic games like in the present experiments can allow researchers to study spending in controlled environments, abstracted from the complexity of real-world spending. The present experiments show how mental accounts and cues of resource scarcity affect spending, while holding aside many factors that would make these effects very difficult to disentangle with observational methods. Future research can use these controlled methods to incorporate and study other key elements from theories of consumer spending.

The present game allows researchers to measure performance, which is usually hidden from the researcher's view in real-world spending. Measuring performance also allows us to characterize participants' cognitive abilities in addition to their inefficiencies. Participants performed substantially better than chance, surviving 40%-60% more days than random consumption across the different variations of the game, which suggests that people have some natural talent for efficient spending. This observation illustrates how performance-based games can help rebalance behavioral economics toward understanding people's good economic decisions as much as their mistakes (Todd & Gigerenzer, 2007).

A performance perspective also fits with evolutionary approaches to behavior in which an animal's decisions to store resources are shaped by their fitness consequences. Previous work emphasizes that non-human animals have trouble delaying gratification and imagining the future (Berns et al., 2007), but in seeming contradiction, a vast array of animal species including

mammals, birds, and insects routinely store food for long periods of time (Pravosudov & Smoulders, 2010; Smith & Reichman, 1984; Vander Wall, 1990). These observations can be reconciled if storage behaviors are shaped not only by general-purpose cognitive abilities to imagine the future (Suddendorf & Corballis, 2007), but also by specialized cognitive mechanisms (De Waal, 2016; Todd & Gigerenzer, 2007; Tooby & Cosmides, 1992) for stockpiling and rationing a reserve of resources.

It is likely that humans also have specialized cognitive abilities for saving and spending reserves, which would explain why across cultures humans store resources for hard times, despite the difficulty of imagining the future and delaying gratification. Even so, people's saving decisions are complex, and the modern economy is much different from the evolutionary past (Seabright, 2004). The economic landscape of efficient saving is considerably altered by modern institutions surrounding money, debt, banks, markets, employment, welfare programs, insurance, and so on. Hence, there are many potential sources of mismatch between our natural saving abilities and modern economic environments. A better understanding of both saving abilities and inefficiencies could ultimately help people meet the distinctive challenges and hardships in modern economic life.

References

- Arkes, H. R., Joyner, C. A., Pezzo, M. V., Nash, J. G., Siegel-Jacobs, K., & Stone, E. (1994). The psychology of windfall gains. *Organizational Behavior and Human Decision Processes*, 59, 331–347.
- Baker, S. R., & Yannelis, C. (2017). Income changes and consumption: Evidence from the 2013 federal government shutdown. *Review of Economic Dynamics*, 23, 99-124.
- Ballinger, T. P., Palumbo, M. G., & Wilcox, N. T. (2003). Precautionary saving and social learning across generations: an experiment. *The Economic Journal*, 113(490), 920-947.
- Barrett, H. C., & Kurzban, R. (2006). Modularity in cognition: framing the debate. *Psychological review*, 113(3), 628.
- Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Analysis*, 20, 351-368.
- Berns, G. S., Laibson, D., & Loewenstein, G. (2007). Intertemporal choice—toward an integrative framework. *Trends in Cognitive Sciences*, 11, 482-488.
- Blake, P. R., & Harris, P. L. (2009). Children's understanding of ownership transfers. *Cognitive Development*, 24, 133–145.
- Blumenthal, J. A. (2010). Property law: A cognitive turn. *Psychonomic Bulletin & Review*, 17, 186–191.
- Boyer, P. (2015). How natural selection shapes conceptual structure: Human intuitions and concepts of ownership. In S. Laurence & E. Margolis (Eds.), *The conceptual mind. New directions in the study of concepts* (pp. 185–200). Cambridge, MA: The MIT Press.
- Brown, A. L., Chua, Z. E., & Camerer, C. F. (2009). Learning and visceral temptation in dynamic saving experiments. *The Quarterly Journal of Economics*, 124(1), 197-231.
- Carbone, E., & Hey, J. D. (2004). The effect of unemployment on consumption: an experimental analysis. *The Economic Journal*, 114(497), 660-683.
- Carroll, C. D. (2001). A theory of the consumption function, with and without liquidity constraints. *The Journal of Economic Perspectives*, 15(3), 23-45.
- Cherry, T. L., Frykblom, P., & Shogren, J. F. (2002). Hardnose the dictator. *The American Economic Review*, 92, 1218-1221.
- DeScioli, P., Karpoff, R., & De Freitas, J. (2017). Ownership dilemmas: The case of finders versus landowners. *Cognitive Science*, 41(S3), 502-522.
- DeScioli, P., & Karpoff, R. (2015). People's judgments about classic property law cases. *Human Nature*, 26, 184–209.
- DeScioli, P., & Wilson, B. J. (2011). The territorial foundations of human property. *Evolution and Human Behavior*, 32, 297-304.
- De Waal, F. (2016). *Are we smart enough to know how smart animals are?* New York: W. W. Norton & Company.
- Frederick, S., Loewenstein, G., & O'donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40, 351-401.
- Friedman, M. (1962). *Capitalism and freedom: With the assistance of Rose D. Friedman*. University of Chicago Press.
- Friedman, O., & Neary, K. R. (2008). Determining who owns what: Do children infer ownership from first possession?. *Cognition*, 107(3), 829-849.
- Gallistel, C. R., Krishan, M., Liu, Y., Miller, R., & Latham, P. E. (2014). The perception of probability. *Psychological Review*, 121, 96-123.

- Ganong, P., & Noel, P. (2017). How does unemployment affect consumer spending? Unpublished manuscript, Harvard University, Cambridge, MA.
- Gelman, M., Kariv, S., Shapiro, M. D., Silverman, D., and Tadelis, S. (2014). Harnessing naturally occurring data to measure the response of spending to income. *Science*, 345(6193):212–215.
- Gross, D. B., & Souleles, N. S. (2002). Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data. *The Quarterly Journal of Economics*, 117(1), 149-185.
- Hoffman, E., McCabe, K., Shachat, K., & Smith, V. (1994). Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior*, 7, 346-380.
- Jappelli, T., & Pistaferri, L. (2010). The consumption response to income changes. *Annual Review of Economics*, 2(1), 479-506.
- Jones, O. D., & Brosnan, S. F. (2008). An evolutionary perspective on the endowment effect. *William and Mary Law Review*, 49, 1935-1990.
- Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1990). Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy*, 98, 1325-1348.
- Kanngiesser, P., Gjersoe, N., & Hood, B. M. (2010). The effect of creative labor on property-ownership transfer by preschool children and adults. *Psychological Science*, 21, 1236–1241.
- Kanngiesser, P., & Hood, B. (2014). Not by labor alone: Considerations for value influence use of the labor rule in ownership transfers. *Cognitive Science*, 38, 353–366.
- Kim, S., & Kalish, C. W. (2009). Children’s ascriptions of property rights with changes of ownership. *Cognitive Development*, 24, 322–336.
- Kokko, H., Lopez-Sepulcre, A., & Morrell, L. J. (2006). From hawks and doves to self-consistent games of territorial behavior. *The American Naturalist*, 167, 901–912.
- Kurzban, R. (2010). *Why everyone (else) is a hypocrite: Evolution and the modular mind*. Princeton University Press.
- Kuijt, I., & Finlayson, B. (2009). Evidence for food storage and predomestication granaries 11,000 years ago in the Jordan Valley. *Proceedings of the National Academy of Sciences*, 106(27), 10966-10970.
- Larrimore, J., Dodini, S., and Thomas, L. (2016). Report on the Economic Well-Being of U.S. Households in 2015. Board of Governors of the Federal Reserve System.
- Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge: Cambridge University Press.
- Milkman, K. L., & Beshears, J. (2009). Mental accounting and small windfalls: Evidence from an online grocer. *Journal of Economic Behavior & Organization*, 71, 384-394.
- Modigliani, F., & Brumberg, R. (1954). Utility analysis and the consumption function. In: Kurihara, K.(Ed.), *Post-Keynesian Economics*. Rutgers University Press, New Brunswick, NJ.
- Morewedge, C. K., & Glibin, C. E. (2015). Explanations of the endowment effect: an integrative review. *Trends in Cognitive Sciences*, 19, 339-348.
- Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer Spending and the Economic Stimulus Payments of 2008. *American Economic Review*, 103(6):2530–53.
- Poterba, J. M. (1988). Are consumers forward looking? Evidence from fiscal experiments. *The American Economic Review*, 78(2), 413-418.

- Pravosudov, V. V., & Smulders, T. V. (2010). Integrating ecology, psychology and neurobiology within a food-hoarding paradigm. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1542), 859-867.
- Prelec, D., & Loewenstein, G. (1998). The red and the black: Mental accounting of savings and debt. *Marketing Science*, 17, 4-28.
- Purves, D., & Lotto, R. B. (2003). *Why we see what we do: An empirical theory of vision*. Sinauer Associates.
- Seabright, P. (2004). *The company of strangers: A natural history of economic life*. Princeton University Press.
- Sherratt, T. N., & Mesterton-Gibbons, M. (2015). The evolution of respect for property. *Journal of Evolutionary Biology*, 28, 1185-1202.
- Smith, A. (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*. Edwin Cannan, ed. 1904. Library of Economics and Liberty.
- Smith, C. C., & Reichman, O. J. (1984). The evolution of food caching by birds and mammals. *Annual Review of Ecology and Systematics*, 15(1), 329-351.
- Souleles, N. S. (1999). The response of household consumption to income tax refunds. *The American Economic Review*, 89(4), 947-958.
- Stake, J. E. (2004). The property 'instinct'. *Philosophical Transactions of the Royal Society of London B*, 359, 1763-1774.
- Suddendorf, T., & Corballis, M. C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral and Brain Sciences*, 30, 299-313.
- Sussman, A. B., & O'Brien, R. L. (2016). Knowing when to spend: Unintended financial consequences of earmarking to encourage savings. *Journal of Marketing Research*, 53(5), 790-803.
- Thaler, R. (1985). Mental accounting and consumer choice. *Marketing Science*, 4, 199-214.
- Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making*, 12, 183-206.
- Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science*, 16, 167-171.
- Tooby, J., & Cosmides, L. (1992). Psychological foundations of culture. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind* (pp. 19-136). New York: Oxford University Press.
- Vander Wall, S. B. (1990). *Food hoarding in animals*. Chicago: University of Chicago Press.
- Wilcox, David W. (1989). Social security benefits, consumption expenditure, and the life cycle hypothesis. *Journal of Political Economy*, 97, 288-304.
- Wilde, P. E., & Ranney, C. K. (2000). The monthly food stamp cycle: Shopping frequency and food intake decisions in an endogenous switching regression framework. *American Journal of Agricultural Economics*, 82(1), 200-213.