OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

The Cost of Resilience

Rom Avissar
romavissar@gmail.com

ABSTRACT

Each month, households that face income uncertainty must decide how to divide what remains after essential spending between savings, debt repayment, and investment. Yet, the advice these households commonly receive comes in the form of fixed rules that do not consider income or market risk. In the paper, the allocation of a household's monthly surplus across savings, debt repayment, and investment is first formulated as a two-stage stochastic program with a CVaR term. Second, it calibrates the model to three archetypes using federal microdata: income level and annual transitions come from the CPS, monthly volatility comes from bank-account data, debt and credit denial come from the SCF, and essential spending comes from the CEX. Third, this paper uses a benchmark comparison on identical income paths to evaluate the optimal allocation policy against the five heuristic policies. In doing so, it measures the expected wealth that the household forgoes in attempting to gain financial security; this is called the cost of resilience. This investigation's central finding is that a buffer held in the form of savings reduces the expected net-worth gain of a household, while getting no additional protection from bad income and market months. Within the optimal policy, the surplus is almost never allocated to savings, and investing the entire surplus is statistically indistinguishable from it. Over two years, three of the four heuristics that hold savings give up between $4,100 and $11,400 of expected net-worth gain. Since greater caution increases the expected cost without reducing CVaR no cost of resilience frontier can be traced.

1 INTRODUCTION

1.1 The Decision Every Young Adult Faces

The median full-year service worker in the United States earns $3,000 per month. At that income, the essential spending for one person is $2,589, leaving a monthly surplus of $411. That same service worker, if indebted, typically owes $1,400 in credit card debt that accrues interest at 21%. In the model they start with no savings or investments. Each of these figures comes from §4, where they are calibrated from U.S. federal microdata.

The $411 can be allocated to three distinct channels: it can be held in savings, the safest option that earns the least interest but faces no risk regardless of the month; it can be allocated to debt repayment beyond the required minimum, reducing the credit card balance that is accruing interest at 21%; or, it can be invested, providing a high market return while also being prone to volatility and risk. Each dollar allocated to one of these channels has an opportunity cost, meaning it cannot be allocated to the other two.

Furthermore, income varies across different types of workers. For example, a salaried professional earns approximately 2.2 times as much as a service worker, providing a monthly surplus that is larger by a factor of eight. However, since essential spending rises far slower than income, for the salaried professional, essentials take a far smaller share of monthly income. That said, for certain archetypes that surplus is not a given; for example, for service workers income falls below the essential floor in 45% of months, meaning that those months leave no surplus at all.

This paper asks: under uncertain monthly income, how should a household divide its limited surplus between emergency savings, debt repayment, and investment, and what does it cost to prioritize resilience over building wealth? That said, households facing such a decision do not model their income, instead they follow simple rules called heuristics.

1.2 Why Folk Heuristics Fail under Income Uncertainty

The modern day financial literacy industry is built around a handful of rules called heuristics: emergency-fund-first, debt avalanche, debt snowball, a fifty-fifty split between savings and debt, equal thirds across savings, debt, and investment, and an invest-first strategy that pays only the required debt minimums. These rules are not irrational. Each was designed around a real financial objective.

Debt avalanche, for example, minimizes total interest paid when future income is known and the household will never need to borrow back money it has already used to repay debt. That is a genuine optimality result. Emergency-fund-first responds to a genuine problem: a household lacking liquid savings may be forced to borrow in the case of a minor, temporary income shortfall. Debt snowball makes the trade-off of costing more in interest, but paying off smaller balances first. This gives households a clearer sense of progress. In behavioral research it is also suggested to help households stay committed to repayment (Gal and McShane). When considering the narrow problems they were built to address each of these rules is defensible.

The main issue with the heuristics is that they all ignore the struggles associated with income uncertainty. For workers paid hourly wages, that uncertainty is very common. The pay they take home varies by about 9% month to month. This occurs even when the worker stays in the same job, and in roughly one month out of four the change exceeds 21% (Ganong et al.). None of these heuristics conditions its decision on the household's income distribution. The same basic advice can therefore be given to a salaried professional and to a gig worker whose income is considerably more volatile. Additionally, the rules are based on the assumption that there will always be a positive monthly surplus. When that monthly surplus turns negative, they implode and offer no guidance. One more assumption they make is that credit will be accessible no matter what. In the model developed in §3, credit can be frozen in months when income falls below the essential floor, so a reduction in income can coincide with an inability to borrow.

The question of whether the heuristics lose to the optimal policy, and by how much, is an empirical one. What matters more is which rules still work under uncertain income, how large the gap becomes, and which households bear the cost.

Instead, this paper models the household's decision directly and solves it under calibrated income uncertainty.

1.3 Contribution

This paper makes three main contributions. First, the allocation of a household's monthly surplus across savings, debt repayment, and investment is formulated as a two-stage stochastic program with a CVaR term. This is then evaluated under a fixed-order recourse that real households could follow as advice; though, this evaluation is bounded by a full-foresight version. Second, it calibrates the model to three archetypes in the workforce using federal microdata: income level and annual transitions come from the CPS, monthly volatility comes from bank-account data, debt and credit denial come from the SCF, and essential spending comes from the CEX. Third, the paper benchmarks the optimal allocation policy that it determines against five commonly recommended heuristic policies on identical income paths. It then measures the wealth that each policy forgoes for financial security, referring to it as the cost of resilience.

The results reject the premise of the common rule of thumb. Holding the monthly surplus in the form of savings is almost never optimal; the heuristics that do so forgo about 8% of the wealth a household could have built over the two years without reducing its unmet need. A cost of resilience frontier cannot be traced because shortfalls are not caused by allocation; instead, they are the result of how much of income essential spending takes, and by access to credit.

2 BACKGROUND

2.1 Precautionary Savings and Liquidity Constraints

Households facing income risk tend to keep more assets in a liquid manner than they would under certainty. Kimball formalized this concept as "prudence," and he showed that convex marginal utility causes risk to increase optimal saving. Deaton added liquidity constraints to this concept. Households with the inability to borrow hold assets as insurance and their consumption tracks income. Even so, Deaton's borrowing constraint was a fixed limit known in advance. Carroll's buffer-stock model shows that an impatient but prudent consumer saves toward a target wealth-to-income ratio and stops saving once that target is reached. Carroll and Summers's research empirically documented that consumption growth parallels income growth for working-age households. This paper's subjects, by contrast, are households actively juggling competing uses of cash month to month.

These models identify a target buffer level and a consumption function within an infinite-horizon setting with one liquid asset. Their results mainly show how optimal behavior changes under different assumptions and derive the corresponding Euler conditions. However, they remain descriptive and do not tell households how to allocate their income each month.

There are several factors limiting these models. First, it is important to consider that they typically contain only one asset. In turn, this means they do not consider the joint allocation decision across liquid savings,

high-interest debt, and risky investment. Second, the models lack an operational policy. The theory explains why households should hold a buffer, while overlooking how to fund that buffer when doing so has an opportunity cost. Third, these models generally consider credit access to be certain. A borrowing constraint may exist, but the household knows from the outset whether it will apply. Thus, the theory establishes the necessity of a buffer, but it overlooks its creation and inaccessibility of credit. In turn, an optimization framework is needed for decisions made prior to the resolution of uncertainty.

Literature that is related to the premise of this paper studies the reason why households hold liquid assets while carrying expensive credit card debt. Gross and Souleles document this joint holding directly in credit card data. Telyukova explains it as the need for cash to cover expenses that cannot be paid by card. On the other hand, Druedahl and Jørgensen explain it through precautionary borrowing against future credit risk. Fulford shows that credit limits are prone to substantial variation, meaning that households cannot rely on a fixed line. These mechanisms give real reasons to hold savings that this paper's model largely excludes: all spending can be met by withdrawals or borrowing, and freezes are rare. Therefore, the finding that savings are not optimal should be read while considering these papers (§6.1).

2.2 Stochastic Programming and CVaR

The standard framework for decision-making under uncertainty is two-stage stochastic programming: decide now, observe, adjust. First, the decision-maker commits a first-stage decision; then, uncertainty realizes; and finally, recourse actions cope with the outcome. The framework incorporates two standard diagnostics. These are the Value of the Stochastic Solution (VSS) and Expected Value of Perfect Information (EVPI). VSS measures the cost of ignoring uncertainty. EVPI measures the value of having perfect foresight.

Even so, it is crucial to acknowledge that risk requires a measure, and the standard measure, variance, fails due to its symmetric nature. No household fears a surprisingly good month, it fears the tail where essentials go unpaid. Value-at-Risk (VaR) identifies a loss threshold. That said, it is limited by its inability to determine the severity of losses beyond it and the fact that it can violate subadditivity. Artzner et al. addressed this problem by defining the coherence conditions that a sound risk measure should satisfy. Conditional Value-at-Risk (CVaR) measures the expected loss within the worst-performing tail of outcomes. It is coherent and directly captures the solvency risk this model is designed to control. Rockafellar and Uryasev showed that CVaR minimization can be written as a linear program using one auxiliary variable. Because of the Rockafellar-Uryasev reformulation, the optimization stays linear.

Generally, stochastic programming is a tool used for institutional asset-liability management. It has applications for pension funds, insurers, and other large financial institutions. A smaller subset of the literature applies this tool to regular households. Consigli, Moriggia and Vitali model a household's long-horizon asset-liability problem as a multistage stochastic program with goals for housing and retirement. Konicz et al. use a combination of stochastic control with multi-stage stochastic programming for personal finance and pensions. Even so, these models do not share this paper's setting. In the case of

their subjects, wealth is accumulated over decades. Then, the same subjects choose between investment vehicles, and plan for liabilities they chose rather than ones forced on them.

In this case, the problem the household faces presents itself through unsecured debt at credit card rates, a lack of assets, the need to meet essential spending every month, and the potential loss of credit access exactly when income falls. This paper's model examines how to split a few hundred dollars of monthly residual across three competing use cases. Even so, the tight binding constraints mean that the smaller scale does not make the problem easier. Precautionary savings theory establishes the reason and need for households to hold buffers. Combining it with stochastic programming creates a method for allocating the monthly surplus under income uncertainty. Common financial advice generally does not combine the two.

2.3 The Gap: No Operational Policy for the Joint Decision

The financial literacy industry is founded on a set of heuristic rules. These include building an emergency fund first, repaying debt using the avalanche or snowball method, dividing income using fixed percentages, or prioritizing investing. Under certain income, the avalanche method is the cheapest because paying the highest-interest debt first minimizes total interest. These rules usually assume that income is known in advance, monthly surplus remains positive each month, and credit is available as a backstop.

Additionally, households receiving financial advice often cannot evaluate it for themselves. Lusardi and Tufano's research explains that debt literacy is low. This shows itself in findings like the fact that many borrowers cannot calculate compound interest on their own balances. This means that the intended audience cannot easily audit the advice it receives, while poor recommendations carry little accountability. Morduch and Schneider's Financial Diaries highlight that variation in income occurs commonly in the modern economic environment. Households regularly run monthly deficits and cope through family loans, juggled bills, and delayed payments. Income variation also differs significantly across different occupations, most reflecting hourly wages versus salaried pay (Ganong et al.). Ganong et al. find that the typical hourly worker's net income changes by about 9% from one month to the next.

This paper joins these three literatures. The monthly allocation is formulated as a two-stage stochastic program with a Conditional Value-at-Risk term in its objective. Its possible allocation channels are emergency fund savings, repaying debt, and investment. Then the model is calibrated using federal microdata across three household archetypes and benchmarks the common heuristics against the optimal policy. This paper measures what those heuristics cost under the uncertainty documented in the Financial Diaries. It also explores their effectiveness when applied by the households they are meant for.

3 MODEL

3.1 Overview

Households with uncertain income have to make the same decision monthly: how to divide the surplus remaining after essential spending between savings, investment, and debt repayment. In this section, that decision will be formulated mathematically as a two-stage stochastic program. In stage one, the household decides fixed shares of its monthly surplus to allocate into one of these; this allocation is decided prior to the arrival of income. In stage two, income is revealed along each simulated path, and the household responds by withdrawing savings, liquidating investments, or borrowing on its credit line to cover any shortfall. The response follows a fixed order and uses only information available that month. A version where the response sees the whole income path is solved as a mixed-integer program and acts as a full-foresight benchmark (§3.2). The objective function minimizes the total expected financial cost plus a penalty incurred on the worst 5% of outcomes, measured using Conditional Value-at-Risk (CVaR). The problem is solved using a sample average approximation over the simulated paths.

The model has two different versions: a multi-period version and a single-period version. The multi-period model tracks balances monthly over a 24-month horizon and produces the main results and conclusions of this investigation. Over a single period, compounding across the months is removed, allowing a demonstration of the model's behavior. Table 1 contains the notation used throughout this investigation.

Symbol	Meaning
Indices	
$t = 1, ..., T$	Month
	Simulated income path (scenario), each with probability $\frac{1}{N}$.
k	Archetype (salaried professional, service worker, gig worker).
y	Year
First-Stage Decisions	
f_S	Shares of monthly surplus allocated to savings, investments, and debt repayment, where $f_S + f_I + f_D = 1$.
f_I	
f_D	

save_0	Initial borrowed amount placed in savings.
invest_0	Initial borrowed amount placed in investments.
borrow_0	Initial borrowed amount, where $\text{borrow}_0 \geq \text{save}_0 + \text{invest}_0$ (equality at every optimum).
Second-Stage Decisions	
w_t	Withdrawal from savings.
v_t	Liquidation of investments.
b_t	New credit line borrowing.
d_t	Debt repayment from surplus, where $d_t = \min(f_D a_t, D_{t-1})$.
m_t	Additional payment needed to reach the minimum debt payment.
u_t	Unmet essential spending.
$u_{m,t}$	Unmet minimum debt payment.
y_t^d	Binary variables enforcing the debt repayment rule and credit limit (bound specification only).
y_t^b	
States	
I_t	Monthly income.
L_y	Annual income in year y.
ϵ_t	Monthly income noise.
a_t	Monthly surplus, where $a_t = \max(0, I_t - E_t)$.
S_t	Savings balance.
V_t	Investment balance.
D_t	Debt balance.

ϕ_t	Credit freeze indicator (1 if credit line is frozen in month t, 0 otherwise).
Parameters	
E_t	Essential monthly spending floor.
r_S	Monthly interest rate on savings.
r_D	Monthly interest rate on debt.
R_t	Realized monthly investment return.
p_m	Monthly probability that the credit line is frozen.
K	Credit limit as a multiple of E_t ($K = 3$).
μ	Minimum payment rate ($\mu = 0.02$).
M	Penalty per dollar of unmet need ($M = 100$).
α	CVaR confidence level ($\alpha = 0.95$).
λ	Weight on the CVaR term.
σ_k	Scale of the monthly income shock for archetype k.
ν	Degrees of freedom of the income shock ($\nu = 0.74$).
P	Annual rank-transition matrix.
$\overline{S_0}$	Initial savings before first-stage decisions.
$\overline{V_0}$	Initial investments before first-stage decisions.
$\overline{D_0}$	Initial debt before first-stage decisions.
Costs and Risk Measure	
C_s	Total cost on path s, including unmet need penalty.
C_s^{fin}	Financial cost on path s (negative when net worth increases).

η	Value-at-Risk threshold.
z_s	Cost on path s in excess of η.
Z	Objective function value (Equation 14).
Z^*	Objective function value under the optimal policy.
Benchmarks	
π	A benchmark heuristic policy.
$Z(\pi)$	Objective function value under heuristic π.
$\delta(\pi)$	Relative cost of π against the optimal policy (Equation 16).
VSS	Value of stochastic solution (Equation 17).
EVPI	Expected value of perfect information (Equation 18).

Table 1 — Notation of the variables used throughout the investigation.

3.2 Timing and Information

In the model, decisions made prior to the arrival of income are distinguished from those made as income arrives. In stage one, the household establishes three shares, f_S, f_I, and f_D, where $f_S + f_I + f_D = 1$. These shares are the percentages of each month's surplus, allocated to savings, investments, and debt repayment. They remain fixed over the 24-month horizon. In turn, this acts as a standing policy, unchanged from month to month. The household is also able to borrow an initial amount, denoted borrow_0. This borrowing can be placed into savings or investments at the start of the 24-month horizon.

In stage two, income is revealed month by month on every simulated path s. In the event where that month's essential spending is not achieved, the household relies on its resources to cover the shortfall. First, the household withdraws its savings (w_t), then it liquidates its investments (v_t), and only finally it borrows, subject to the credit limit placed on it, using its credit-line (b_t). Any shortfall remaining after the three resources are used is considered to be unmet need, denoted u_t or $u_{m,t}$, and penalized in the objective function. The sequence presented is based on standard financial advice and relies only on the month t's available information. The same recourse is followed by the optimal policy and every benchmark heuristic. This means their expected cost differences result only from allocation, including the optimal policy's option to borrow at month 0, which the heuristics do not use.

To establish a full-foresight benchmark for the best possible expected cost, stage two is also solved as a mixed-integer program that can access the entire income path for the 24 months. This allows it to select all recourse actions jointly. That said, no real household could use this policy. Therefore, it serves as a full-foresight benchmark; it would be a lower bound if every search converged, while the difference between the two specifications bounds the value of foresight (§5.6). The true optimum, over all policies including those whose shares adapt over time, lies somewhere between them. In Figure 1, the timings within a month are summarized.

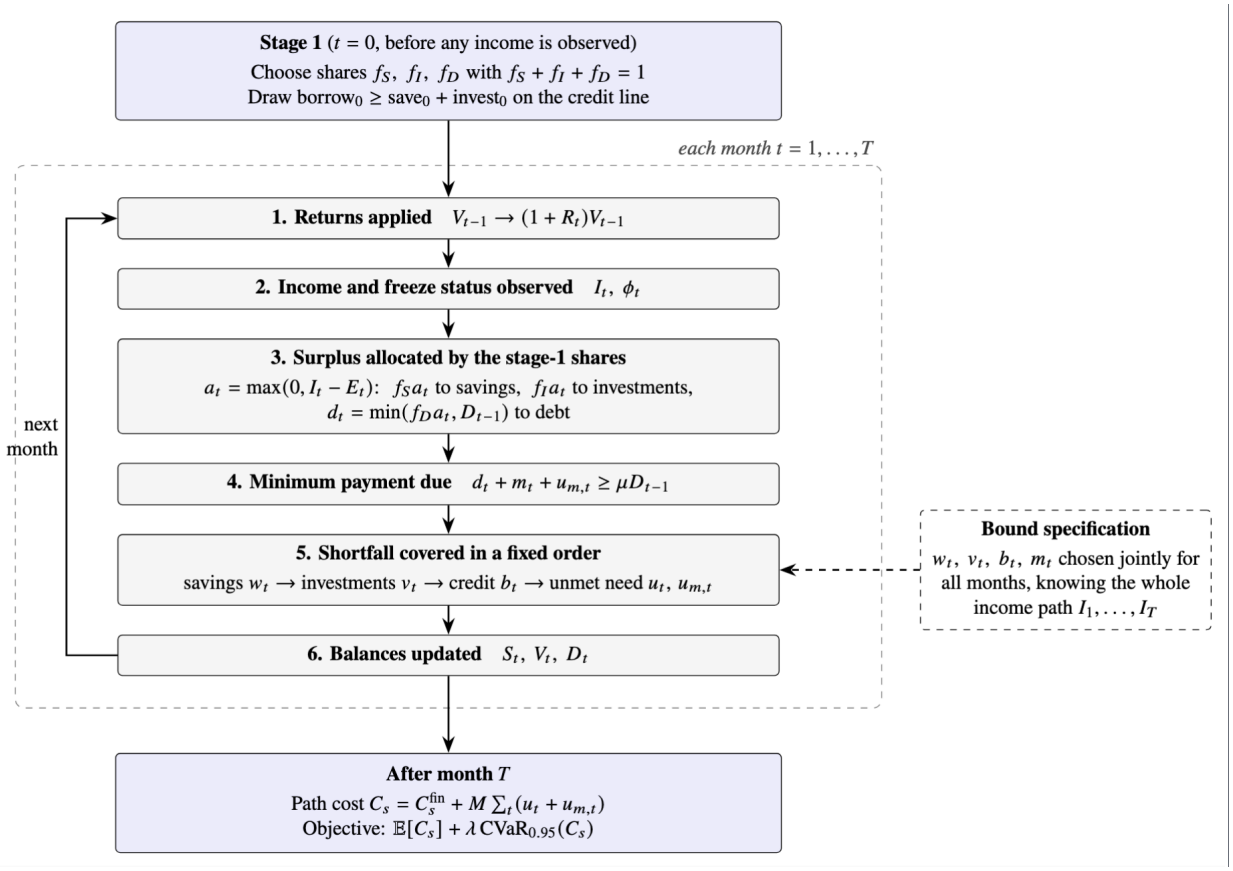


Figure 1 — Timing of decisions within a month.

3.3 Income Process

Monthly income is modelled as the product of two components, a persistent annual level and a transitory monthly shock:

$$I_t = \frac{L_{y(t)}}{12}(1 + \epsilon_t). \tag{1}$$

where $L_{y(t)}$ represents the household's annual income in year y, and ϵ_t is the deviation of month t from the year's average. Here, $\frac{L_{y(t)}}{12}$ expresses the changes in a household's earnings from one year to another. $(1 + \epsilon_t)$ captures the monthly volatility described in the Financial Diaries and bank-account data.

The annual income level, $L_{y(t)}$, is persistent. Each household has a rank within its archetype's annual income distribution. Over the years the household's rank changes according to a transition matrix estimated from consecutive-year observations of the same workers. $L_{y(t)}$ is drawn from the incomes of full-year workers in the household's current quintile. Once a year, the transition occurs, in a month drawn at random for every path. Therefore, annual changes in a household's income do not all occur in the same calendar month across paths.

Each month, ϵ_t is drawn independently from a Student-t distribution:

$$\epsilon_t \sim t_\nu(0, \sigma_k^2) \text{ truncated to } [-0.9, 0.9]. \tag{2}$$

where $t_\nu(0, \sigma_k^2)$ denotes a Student-t distribution with location 0, scale σ_k, and ν degrees of freedom. The truncation keeps income positive, and scale σ_k differs by archetype. The shape of this distribution captures the rare but large income swings that a normal distribution, with thinner tails, would understate. Because the distribution is symmetric, ϵ_t has a mean of zero, which means that a shock shifts income between months without changing the expected annual level. The transition matrix, degrees of freedom, and scale of ϵ_t are estimated in §4.

3.4 Cash Flow and Debt Rules

Each month begins with the household's surplus, the income left after essential spending:

$$a_t = \max(0, I_t - E_t). \tag{3}$$

The shares in stage one divide this surplus. The household places $f_S a_t$ in savings and $f_I a_t$ in investments, then it pays

$$d_t = \min(f_D a_t, D_{t-1}) \tag{4}$$

towards its debt, so it can never repay more than it owes. Months where income falls below the floor mean that $a_t = 0$ and nothing is allocated to any of the three possible options.

To prevent delinquency the household must also perform a minimum payment towards its debt. If the planned repayment d_t falls short of the minimum it must also perform an additional payment m_t. Any amount that can still not be paid is recorded as $u_{m,t}$:

$$d_t + m_t + u_{m,t} \geq \mu D_{t-1}, \tag{5}$$

where $\mu = 0.02$. All spending in month *t* must then be funded:

$$I_t + w_t + v_t + b_t + u_t \geq E_t + a_t + m_t, \text{ where } 0 \leq u_t \leq (E_t - I_t)^+ \tag{6}$$

Here, the left side lists every source of cash: income, withdrawals from savings, liquidations of investments, new borrowing, and unmet need. The right side lists every use: essential spending, the allocated surplus, and any top-up to the minimum payment. Due to the bound on u_t, unmet need can only cover a genuine shortfall in essential spending. This also means that it cannot be used to avoid spending the surplus. The same constraint as a monthly flow of funds is shown in Figure 2.

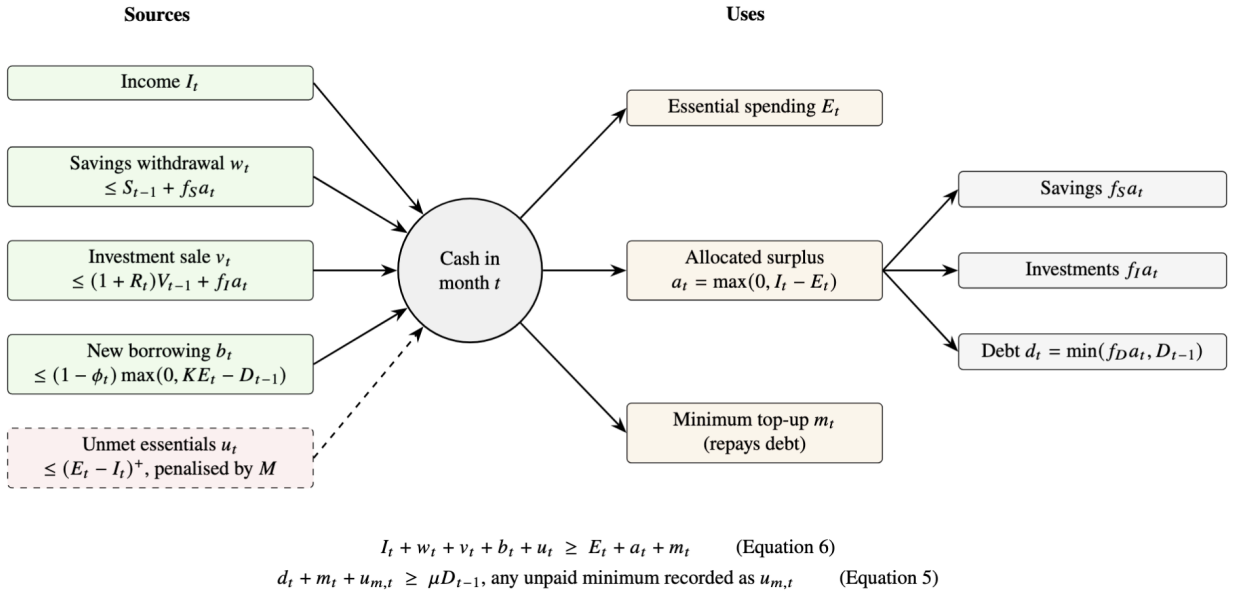


Figure 2 — Cash Flow.

The previous balance plus the current month's contribution limit withdrawals and sales. The credit line limits new borrowing:

$$w_t \le S_{t-1} + f_S a_t, \; v_t \le (1 + R_t) V_{t-1} + f_I a_t, \\ b_t \le (1 - \phi_t) \cdot \max(0, KE_t - D_{t-1}). \tag{7}$$

K, which is equal to three, is a multiplier that acts on essential spending. The product, less debt already owed, is the credit limit. When the credit line is frozen the freeze indicator ϕ_t equals one. Freezes are assigned to the lowest-income months below the floor, until a share p_m of all months are frozen. In practice, they fall only in months of near-zero income and affect about 1% of paths.

As presented in the main specification, the recourse allocations follow a fixed order. A shortfall is only recorded as unmet need after savings, investments, and borrowing up to the credit line are all exhausted. In the bound specification, w_t, v_t, b_t, and m_t are chosen jointly across the whole path by the mixed-integer program. Because neither the minimum in Equation 4 nor the maximum in Equation 7 can be written as a linear constraint, each is enforced by the binary variables.

3.5 Balance Dynamics

The decisions in §3.2 and §3.4 determine how three main balances evolve each month. The savings balance evolves by earning interest at a fixed monthly rate, gaining the money allocated by the household, and losing any withdrawal done by the household:

$$S_t = (1 + r_S) S_{t-1} + f_S a_t - w_t. \tag{8}$$

Investments evolve in a similar way, but instead of growing at a fixed monthly rate they change by a realized monthly return R_t:

$$V_t = (1 + R_t) V_{t-1} + f_I a_t - v_t. \tag{9}$$

Debt evolves similarly, it accrues interest at a fixed rate, but a payment causes it to fall, whereas new borrowing increases it:

$$D_t = (1 + r_D) D_{t-1} - d_t - m_t + b_t. \tag{10}$$

All three balances are non-negative values. At the start of the horizon, the household's initial borrowing and allocations enter each balance:

$$S_0 = \overline{S_0} + \text{save}_0\,,\; V_0 = \overline{V_0} + \text{invest}_0\,,\; D_0 = \overline{D_0} + \text{borrow}_0\,, \tag{11}$$

where the barred terms are the household's holdings prior to any stage one decisions being made. Within a month, investment returns are applied first, then the surplus is allocated, then any shortfall is covered. Withdrawals can also draw on that month's contribution.

The use of realized rather than expected returns is crucial to this investigation. With expected returns, investments would grow smoothly and could always be sold at a predictable value, which would make them strictly better than savings on any possible path. With realized returns, a household that must sell investments after a fall in the market sells at a loss. Even so, no taxes or transaction fees are present in the model, which means that investments are still treated slightly more favorably than they would in reality. The returns, R_t, are drawn from historical data, which will be explored further in §4.

3.6 Objective Function

The model measures outcome by the change in a household's net worth over the horizon. Financial cost is the negative of that change:

$$C_s^{fin} = -[(S_T + V_T - D_T) - (S_0 + V_0 - D_0)]\,. \tag{12}$$

Therefore, a household that ends with a negative C_s^{fin} value has had an increase in net worth, or a negative cost (gain). Any gain or loss in net worth is the result of surplus allocated to savings and investment, borrowing used to cover shortfalls, interest earned on savings, realized investment returns, repayment of debt, and interest paid on debt.

Afterwards, unmet need is added to cost as a penalty:

$$C_s = C_s^{fin} + M\sum_{t=1}^{T}(u_t + u_{m,t}) \quad,\quad M = 100\,. \tag{13}$$

Through the penalty each dollar of missed essential spending or missed minimum debt payment becomes one hundred times costlier than a dollar of lost wealth. This reflects the assumption that missing rent or defaulting on a payment is significantly worse than its equivalent loss in net worth. Such outcomes can in the worst case scenario lead to consequences like eviction or default, which can leave a prolonged strain on a household's economic sustainability and ability to function regularly. Hence, C_s represents well-being rather than a mere dollar amount; C_s is reported in §5 and it should be read in the appropriate units.

The household chooses its first-stage policy to minimize expected cost plus a penalty on the tail:

$$\min \frac{1}{N} \sum_{s=1}^{N} C_s + \lambda \mathrm{CVaR}_{\alpha}(C) \quad , \tag{14}$$

where CVaR_{α} is the average cost over the worst $(1-\alpha)$ share of paths. With $\alpha = 0.95$, it averages the worst 5% of paths. Following Rockafellar and Uryasev, CVaR can be written as:

$$\mathrm{CVaR}_{\alpha}(C) = \min_{\eta} \eta + \frac{1}{(1-\alpha)N} \sum_{s=1}^{N} z_s \quad , \; z_s \geq C_s - \eta, \; z_s \geq 0 . \tag{15}$$

This in turn ensures the bound specification stays linear apart from its binary variables. In the main specification, CVaR is calculated directly from the simulated costs.

The reason CVaR is used instead of variance is because variance would penalize good and bad outcomes equally. In reality, no household fears a month that is unusually successful compared to what it expected. CVaR only penalizes the worst outcomes, which in this model are paths where a household cannot meet its essential spending or minimum debt payment. The importance of the tail is set by λ, where $\lambda = 0$ represents a household that only cares about the expected cost, and higher values represent households that become increasingly more cautious.

3.7 Single-Period Variant

The single-period version maintains the structure of the main model but it removes time's effect on net worth. The household chooses which policy it will use for the period during the first stage. Its income is then drawn for a single month on every path. The household responds to any shortfall using the same sources of cash, limits, and penalties discussed in §3.4 and §3.6. Since there are no future months in this version there is no compounding across balances. Hence, the minimum and maximum in Equations 4 and 7 can be written with linear constraints alone. Therefore, this version of the model is a pure linear program, and a solution is proven optimal by the solver.

This version also removes the distinction between the two recourse specifications. Only one period means that the household observes all relevant income across the model's horizon prior to any recourse actions, so the household knows all relevant income when it acts, meaning foresight does not play a role. The linear program chooses the cheapest source of cash rather than following a fixed order. The accumulation of savings and investments, the compounding of debt, and the deterioration of a household's buffer are not considered by this version because it only runs over a single month, meaning those factors cannot be

accounted for. In turn, this version of the model is only used to highlight how the model trades expected cost against tail risk; no conclusions in §5 are solely based on it.

3.8 Benchmarks and Diagnostics

To determine its relative performance, the optimal policy is benchmarked against rules based on common advice. Unlike the optimal policy, these rules do not have fixed shares, instead, they depend on the current balances of a household.

Policy	Allocation Rules
Emergency Fund First (EFF)	The entire surplus is allocated to savings until three months of essential spending are achieved. Then, everything is directed to debt repayment. Once debt is clear it is allocated to savings again.
Debt Avalanche (DA)	The entire surplus is allocated to debt repayment until D_t is cleared. Then, the entire surplus is directed to savings; this is the paper's implementation.
Invest First (IF)	The entire surplus is allocated to investments.
Fixed Split 50/50 (FS50)	The surplus is allocated equally between savings (50%) and debt repayment (50%).
Fixed Split 33/33/33 (FS33)	The surplus is allocated equally between savings (33%), investment (33%), and debt repayment (33%).

Table 2 — Heuristics used for benchmarking in the investigation.

The same income paths, credit limit, minimum payment, and cost function apply to all policies, including the optimal policy. All of them cover shortfalls using the same recourse (§3.2). The policies' evaluation will be done using two different payment orders: one where the minimum debt payment is paid before essential spending, and one where essentials are paid first.

For any heuristic π, the relative cost is measured by its divergence from the optimal policy:

$$\delta(\pi) = \frac{Z(\pi) - Z^*}{|Z^*|}, \tag{16}$$

where Z denotes the value of the objective function in Equation 14 and Z^* is its value under the optimal policy.

Two further measures assess the value of modelling uncertainty. The value of the stochastic solution (VSS) compares the optimal policy with the expected-value policy, which is chosen on a single path of average income and then evaluated on every simulated path:

$$\mathrm{VSS} = Z(\text{EV policy}) - Z^{*}. \tag{17}$$

The expected value of perfect information (EVPI) compares the optimal policy with one that chooses its first-stage decisions separately for each path, knowing that path's income in advance:

$$\mathrm{EVPI} = Z^{*} - Z(\text{perfect information}). \tag{18}$$

The cost of assuming that its next income will be the average, and ignoring uncertainty is measured by VSS. The highest worth that better information would have for the household is measured by EVPI.

3.9 Solution Method

A sample average approximation is used to solve both specifications. Each candidate policy is scored on a set of 200 simulated income paths. Then, the most successful policy is re-evaluated on a separate set of 2,000 paths; this ensures that the reported costs are not biased upwards by evaluating a policy on the same paths used to choose it. The whole procedure is repeated with five random seeds, and results are reported with confidence intervals across seeds.

The stage one policy is selected using a search in the main specification. In the global grid shares are varied in steps of 0.05. The initial borrowing is varied in ten steps up to the credit limit, and the division of that borrowing between savings and investments in steps of 0.25. A local search then refines the best grid point in steps of 0.01. For each value of λ, the optimal policy is taken from the same set of evaluations. This is done since every candidate's expected cost and CVaR are stored. Since the model is using a numerical search, the policy cannot be considered a mathematically proven optimum. Thus, this investigation considers it optimal within the grid resolution.

In the bound specification, each path's recourse problem is a mixed-integer program, solved with HiGHS (version 1.15.1; Huangfu and Hall) through PuLP (Mitchell et al.). This version of HiGHS occasionally returned suboptimal solutions while reporting them as optimal. To account for this, each problem was solved twice: first, with presolve disabled; then, with a fixed random seed. The lower cost of the two is kept. The two settings of HiGHS occasionally disagreed, in which case the lower proven cost was kept. The retained HiGHS solution on 75 of 75 final-model problems is also supported by CBC. However, on its default settings HiGHS returned one solution that was $12,997 too high. Additionally, it is important to note that 13 of the 60 bound searches reached their evaluation limit; this means that the full-foresight benchmark may be slightly loose. The single-period model is a linear program and is solved directly.

3.10 Assumptions

The model presented in this section rests on several simplifying assumptions. Each assumption is stated below in Table 3 with the direction in which it is likely to bias the results.

Assumption	Explanation
Fixed Shares	Allocations are fixed all throughout. This understates reality, where households are able to adapt, which would create more effective allocations. The true optimum lies below the cost of the main specification.
Recourse	Shortfalls are covered in a fixed order in the main specification, without foresight. In the bound they are covered with the addition of foresight. When combined they provide an upper and lower bound for the cost of the best possible recourse.
No Other Ways to Cope	In reality, households can receive help from family or friends, cut spending below their essential floor by using substitute goods, or borrow through different channels. Unmet need is likely lower, and should be read as an upper bound.
Frictionless Investment	Investments can be sold at any time without any taxes or transaction fees. In reality, this is not the case, however in the model it provides a slight unfair advantage to investments over savings. Any result that states savings are not optimal should be read with this in mind.
Fixed Rates	Interest rates for savings and debt are constant within the model, a minor simplification on a 24-month horizon.
A Single Debt Balance	Only one type of debt is held causing Debt Avalanche and Debt Snowball to become the same policy. In reality, households could also borrow through different channels at different interest rates.
Exogenous Freezes	Credit freezes occur at a fixed monthly rate in low-income months. They are not triggered by the household's own borrowing or missed payments. Credit risk plays a small role in the results because freezes fall only in months of near-zero income.
One Representative Household	I_t & L_y vary across paths, covering the earnings spread within an archetype. $\overline{D_0}$, E_t, r_D, r_S, and allocation do not. In reality,

	those also vary within an archetype.
No Discretionary Spending	In reality, households do not only spend money on their essentials. In the model, there is no additional spending above the essential floor, which makes it possible to allocate its entire surplus into one of the three options. In turn, wealth gains are optimistic and unmet need is understated.
Single-Person Household	Each household represents one adult, one income, and one person's essential floor. Thus, families or households with multiple people are not modelled. A second earner would increase income and reduce unmet need, while a dependent would raise the essential spending floor and unmet need.

Table 3 — Simplifying assumptions present in the model.

4 DATA AND CALIBRATION

4.1 Income Distribution and Annual Transitions

Income is calibrated from the Current Population Survey Annual Social and Economic Supplement (CPS ASEC), accessed through IPUMS (Flood et al.), for income years 2018 to 2024. Each archetype is defined by class of worker and occupation. ASEC person weights are used for all statistics regarding these archetypes.

Archetype	Definition
Salaried Professionals (SPA)	These are wage earners in professional and managerial occupations.
Gig Workers (GWA)	These are unincorporated self-employed workers in any occupation. 29.6% of them work in service and sales occupations. Thus, rather than representing a specific platform of work, the archetype represents self-employment broadly.
Service Workers (SWA)	These are wage earners in service and sales occupations. Sales make up 36.0% of the weighted sample.

Table 4 — Archetype definitions used for this investigation.

To be part of the sample someone has to work at least 50 weeks in the previous year, which is considered to be a full work year. To get monthly income, annual income is divided by twelve. This restriction ensures that workers who worked part of the year and earn accordingly do not skew the results. For example, if someone worked six months in the previous year they would appear to earn half of their true monthly income. By including this, this investigation can prevent blurring the distinction between low pay and time out of work. As a result, 15%, 27%, and 26% of SPA, GWA, and SWA samples respectively are removed from the pool. The median monthly income also rises by 7%, 15%, and 20%. Table 5 reports the results for each archetype.

Archetype	n	Median Monthly Income	Mean Monthly Income	Linked Pairs
SPA	139,769	$6,671	$8,931	34,788
GWA	21,369	$3,350	$5,456	4,735
SWA	85,535	$3,000	$4,484	17,589

Table 5 — Full-year CPS samples by archetype (weighted).

Here, n is unweighted, medians and means are weighted. Linked pairs refers to consecutive year-observations of the same full-year worker.

The full-year workers within each archetype are split into five weighted quintiles. These quintiles are used to form the states of the annual income process. The CPS panel is used to estimate the transitions between states as households surveyed in consecutive years are linked, giving the pairs of annual income in Table 5. The transition matrix P records the share of workers in each quintile in one year who are in each quintile the next, with each year classified by its own quintile cut-offs. Because income is classified by rank rather than by dollar amounts, the matrix ensures that ordinary wage growth does not move households into higher states and simulated incomes are not inflated.

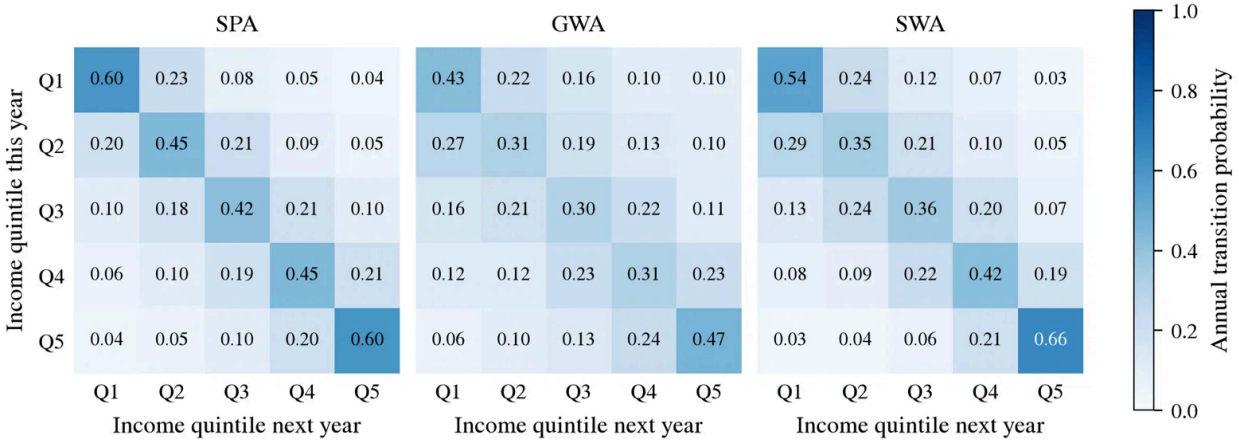


Figure 3 — Estimated annual transition probabilities for each archetype.

Figure 3 shows the estimated transition matrices for each archetype. Income tends to be more persistent at the extremes: for the top quintile SPA has a 60% chance of staying there the following year, while SWA has a 66% chance of doing so; for the bottom quintile SPA has a 60% chance of staying there the following year, while SWA has a 54% chance of doing so. GWA incomes are the least persistent, with diagonal probabilities between 30% and 47%.

Each path begins in a state drawn from the stationary distribution of P. Whenever the household enters a state, its annual income L_y is drawn, using the survey weights, from the income of full-year workers in that quintile. As described in §3.3, transitions occur every twelve months, starting in a month drawn at random for each path. Because the bottom quintile of self-employed workers includes some with no earnings in a given year, 1.3% of simulated GWA months have zero income.

The CPS distribution is almost identically reproduced by the calibrated process. For a single year's income, the ratio of simulated median annual income to CPS median annual income is: 1.003 for SPA,

1.033 for GWA, and 0.996 for SWA. For all three archetypes the ratio means lie between 0.99 and 1.00. Measured over fixed calendar years, the median ratios become 1.027, 1.091, and 1.015. These are increased because mid-year transitions combine two different income levels within the calendar year. Even so, in the multi-year measure, the only archetype to depart from data by more than 3.0% is GWA, at 9.1%. Thus, in this investigation GWA households are slightly better off than in reality.

It should be noted that this sample includes the pandemic years, which may have had unusual transitions. If the investigation were to exclude the links whose income year is 2020 and 2021, it would remove 45% to 48% of links but change no transition probability by more than 0.031; it would also change none of the conclusions (§5.10). In turn, the impact of keeping these years is not detrimental to the validity of this investigation. Additionally, it is important to emphasize that each archetype describes a single earner, consistent with the one-person essential floor.

4.2 Monthly Income Volatility

The monthly shock ϵ_t is calibrated to bank-account data; so, because the CPS reports income only once a year. Ganong et al. used the accounts of JPMorganChase customers to measure how the take-home pay of hourly workers changes from month to month. Through this they found that the typical change is about 9% and roughly one in four months sees a change of 21% or more. These figures are effective measures of the volatility expected by a household because they follow the same workers over time. Thus, for volatility, the investigation relies on them because the cross-sectional CPS cannot produce the results that are as accurate.

Therefore, these figures justify using the Student-t distribution displayed in Equation 2. Its heavier tails mean that it is possible to produce a small typical change alongside frequent large ones, representing the 9% typical change, while allowing for the 21% figure to also be accounted for in a similar frequency to the one observed in the data. The distribution's parameters are set jointly for SWA, the archetype most similar to the hourly workers presented in the bank-account data. To ensure this, the scale σ_{SWA} and the degrees of freedom ν are adjusted to represent the fixed annual level within a year. Their adjustment is not finalized until the simulated paths reproduce a median absolute month-to-month change of 9% and a 25% share of changes above 21%. The result of this is $\nu = 0.74$ and $\sigma_{\text{SWA}} = 0.0426$. It is important to note that in this case σ_k represents the scale of the distribution rather than a standard deviation. The heavy-tailed nature of monthly pay is reflected by the low value of ν. That said, it requires the truncation displayed in Equation 2, as with ν below one, the untruncated distribution would have no finite mean. Therefore, draws are taken from the exact truncated distribution, which is symmetric, so ϵ_t has a mean of zero.

Because comparable data is not available for the other two archetypes their scales are set relative to SWA. For GWA the scale is set to $\sigma_{\text{GWA}} = 1.5\sigma_{\text{SWA}}$, reflecting the greater volatility of self-employment income. For SPA the scale is set to $\sigma_{\text{SPA}} = 0.5\sigma_{\text{SWA}}$, reflecting more stable pay.

These relative values are assumed. They are not estimated from data. In Table 6, the resulting volatility is reported.

Archetype	σ_k	Median Monthly Change	Share of Changes Above 21%
SPA	0.0213	4.9%	15.2%
GWA	0.0640	12.5%	33.1%
SWA	0.0426	9.0%	25.1%

Table 6 — Calibrated monthly income volatility.

Here, changes are measured within a year at a fixed annual level. The degrees of freedom ν are equal to 0.74 for all archetypes. The inclusion of months with an annual transition raises each median by about one percentage point and each tail share by about four.

Varying σ_{SPA} between 0.25 and 1.00 times σ_{SWA}, or σ_{GWA} between 1.0 and 2.0 times σ_{SWA}, changes the objective by less than 0.3% (§5.10). The tail of outcomes is driven by low annual income levels and months of near-zero income, not by month-to-month noise.

When done simultaneously with the annual income process described in §4.1, this volatility determines how often income falls below the essential floor. In the simulated paths 15% of SPA, 41% of GWA, and 45% of SWA months fall below the essential floor (§4.4).

4.3 Debt, Credit, and Interest Rates

The Federal Reserve's 2022 Survey of Consumer Finances (SCF) is used to calibrate debt (Board of Governors, *Survey of Consumer Finances*). SCF records the credit card balance each household carries after its last payment. That said, only credit card debt is modelled, meaning other types of debt like student loans, auto loans, and mortgages are excluded. The reason for this is because such forms of debt have lower rates and fixed repayment schedules that must be followed; in turn, they do not directly compete with savings and investments in the same way that credit card debt does.

For each archetype, debt is measured among SCF households with incomes that fall within 20% of that archetype's median annual income from §4.1. For example, for SWA those bounds would result in annual incomes of $28,800 to $43,200. Table 7 shows the median incomes and the share of households that have no credit card debt for each of these archetypes.

Archetype	Median Annual Income	Share of Households Without Credit Card Debt
SPA	$80,053	42.6%
GWA	$40,200	52.8%
SWA	$36,000	49.6%

Table 7 — Median annual income and share of households that have no credit card debt.

This investigation and its model target specifically the households that do carry credit card debt. Therefore, the initial debt balance $\overline{D_0}$ is set to the weighted median balance among indebted households in each window: $2,500 for SPA, $1,600 for GWA, and $1,400 for SWA. The number of households that produce each of these values of $\overline{D_0}$ is: 327 for SPA, 243 for GWA, and 204 for SWA.

This investigation uses the median because a few extremely indebted households raise the mean level of indebtedness far beyond the typical household's. In the same stated windows the mean balance among the indebted is: $5,836 for SPA, $3,767 for GWA, and $3,312 for SWA. These figures are each more than twice the values of their respective medians.

SCF incomes are measured for an entire household in the income year 2021. By contrast, the CPS figures are measured for a single-worker's earnings between the income years 2018 and 2024. This distinction therefore causes debt to become an approximation. Additionally, households across the archetypes start with no savings or investments, meaning that $\overline{S_0} = \overline{V_0} = 0$. Any cash at the start must come from the month-0 draw on credit. The Federal Reserve's commercial bank interest rate on credit card plans for the first quarter of 2026 (FRED series TERMCBCCALLNS; Board of Governors, "Commercial Bank") is used to determine that debt accrues interest at $r_D = 21.00\%$ annually. The June 2026 average of 30-day Secured Overnight Financing Rate (SOFR30DAYAVG; Federal Reserve Bank of New York) states that the annual interest rate earned on savings is $r_S = 3.61\%$. However, this approximates the yield on a money-market fund or high-yield savings account, and thus is above the average savings-account rate. In turn, savings are treated generously under this calibration method; therefore, any findings that savings are not optimal are reinforced. The compounding formula is used on both of these annual rates to convert them to monthly rates:

$$r_{\text{monthly}} = (1 + r)^{\frac{1}{12}} - 1, \tag{19}$$

and held fixed within each run. The minimum debt repayment rate is $\mu = 0.02$, which is an assumption of this model.

The credit limit for a household is set at three months of its essential spending, also given by $K = 3$. The resultant credit limits for each archetype are \$9,812 for SPA, \$7,824 for GWA, and \$7,768 for SWA. That said, this limit is only set on new borrowing that happens beyond $\overline{D_0}$. Interest can push debt above the limit because a balance at the limit compounds at a monthly nominal interest rate of 1.6% while the limit grows with inflation at 0.34% per month. Under the optimal policy, this happens on 11% of SPA paths, 37% of GWA paths, and 40% of SWA paths. The most that the debt balance goes over the limit is 30%. A household in this position cannot borrow again until its balance falls back below the limit. Since SCF does not report credit limits in the necessary format K is assumed.

4.4 Essential Spending and Credit Freezes

The essential floor at the beginning of the horizon E_0 is calibrated from the Bureau of Labor Statistics Consumer Expenditure Survey (CEX) for 2023 and 2024, using its table for one-person consumer units by income (Table 3404; U.S. Bureau of Labor Statistics, "Table 3404") to match the single-earner households in the model. Essential spending is calculated as a sum of seven crucial categories: food at home, housing (including utilities, household operations, and furnishing), apparel and services, healthcare, gasoline (and other fuels), other vehicle expenses, and public transportation. Since minimum debt repayment is modelled separately, credit card payments are excluded (§3.4). In the CEX report, spending is stated by income bracket. Therefore, each archetype's floor is interpolated linearly between the brackets on either side of its full-year median income from §4.1. This gives monthly floors of \$3,271 for SPA, \$2,608 for GWA, and \$2,589 for SWA.

Essential spending rises much more slowly than income. Between the CEX brackets with mean incomes of \$34,615 and \$82,119, income rises by a factor of 2.4, while essential spending rises only by a factor of 1.3. Therefore, different shares of each archetype's income are taken as the essential floor. In turn, the floor becomes 49% of SPA median income, 78% of GWA median income, and 86% of SWA median income. For the median archetype, the monthly surplus is: \$3,400 for SPA, \$742 for GWA, and \$411 for SWA. Thus, building the recommended (three months of essentials) fund (Scott) would take 2.9 months for SPA, 10.5 months for GWA, and 18.9 months for SWA. A large portion of workers across the archetypes in the CPS sample fall short of the necessary floor: 12% of SPA, 39% of GWA, and 42% of SWA. In the simulated paths, 15%, 41%, and 45% of months respectively fall below it.

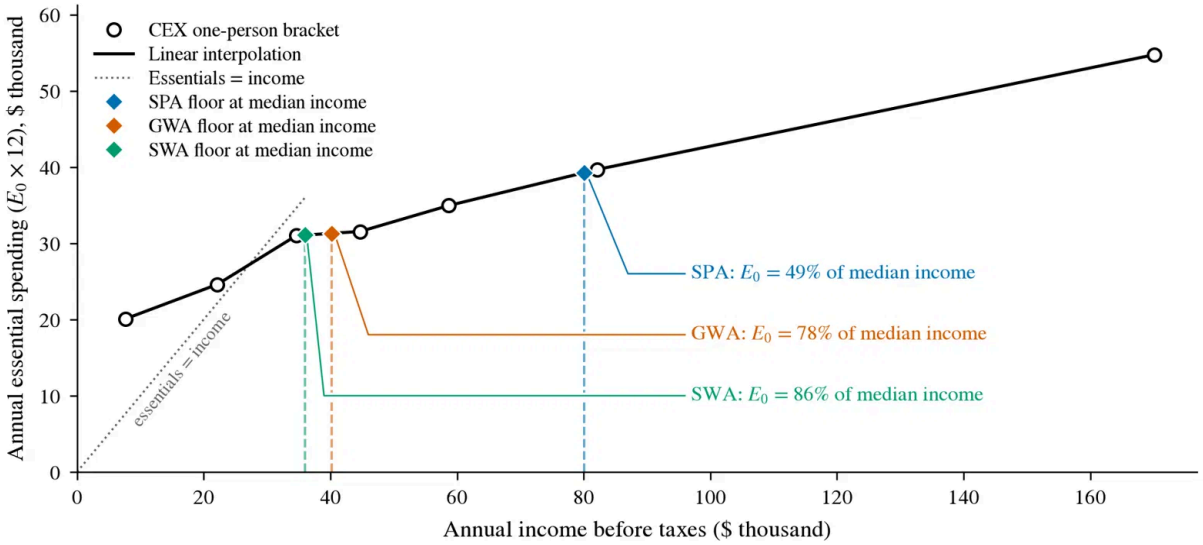


Figure 4 — CEX Income vs. Essential Spending.

Because of inflation the essential floor increases at an annual rate of 4.17%. This figure comes from the change in CPI-U between May 2025 and May 2026 (FRED series CPIAUCSL; U.S. Bureau of Labor Statistics, "Consumer Price Index"). This gives an equation for essential spending in month t:

$$E_t = E_0(1+\pi)^t. \tag{20}$$

By month 24 ($t = 23$), the final month of the multi-period model, essential spending is 8.1% higher than it was at the beginning of the horizon. Income on the other hand, is nominal (as given by the CPS data), which means that it is not adjusted for inflation and does not experience wage growth. The model therefore assumes that real incomes fall over the horizon. This adds about 1.5 to 1.9 percentage points to the share of months below the floor; with the floor held at E_0, the shares would be 14%, 39%, and 43% respectively.

Credit freezes are calibrated from the SCF question on whether a household has been refused credit in the past five years (Board of Governors, Survey of Consumer Finances; variable TURNDOWN). The SCF income quintile closest to each archetype's income gives this rate, resulting in: 12.19% for SPA, 13.16% for GWA, and 12.82% for SWA. These given five-year rates are transformed into monthly rates using the equation:

$$p_m = 1 - (1 - p_{5\text{yr}})^{\frac{1}{60}}, \tag{21}$$

giving monthly rates of 0.22% for SPA, 0.23% for GWA, and 0.23% for SWA. Freezes are assigned to the lowest-income months below the floor until a share p_m of all months is frozen. In practice, this means that within this investigation freezes only occur in months where income nears zero, which means that only about 1% of paths ever encounter a freeze.

4.5 Investment Returns

Investments in this investigation are modelled as a broad U.S. stock index fund. Monthly returns R_t come from Kenneth French's data library. They are the nominal total return on the value-weighted U.S. stock market, computed as Mkt-RF plus RF, for the 240 months from June 2006 to May 2026. It is important to consider that within this window both the 2008 financial crisis and the 2020 pandemic crash are present. On each simulated path, the returns are drawn independently with an equal probability across the 240 months. In turn, paths include both crashes and recoveries that occurred historically, meaning that returns are not as stable as the fixed rate for savings r_S. Policies are evaluated on the same draws, meaning there is no return bias across the results.

Over the sample, the average monthly return is 1.01%, or 12.8% a year when compounded, with a monthly standard deviation of 4.5%. The realized compound return over the period was lower, at 11.5% per year. In the sample, the worst month is October 2008 (the financial crash), at -17.1%; the worst twelve months occur between March 2008 and February 2009, at -42.6%.

This design does result in two important simplifications. First, it is important to note that because returns are drawn independently there is no clustering of bad months that would typically occur in economic downturns, like they did in 2008. While repeated bad months are possible, they are the result of chance. Second, it is important to note that returns are drawn independently from income. In reality, a recession causes income to drop alongside returns, because a market downturn causes unemployment, or more rarely due to their sticky nature, wage cuts. Therefore, in reality, a household is more likely to have to sell its investments at a loss to prevent insolvency. These two simplifications cause investments to be treated more favorably than they would in reality (§3.10).

The high historical average also favors investing, so §5.10 lowers the mean returns to 7.0% a year by shifting every monthly return down by the same amount, leaving their volatility unchanged.

5 RESULTS

5.1 Single-Period Behavior

In the single-period version, the model's behavior and choice between expected cost and tail risk is made evident. Across λ from 0 to 10, the optimal share of surplus placed in savings is zero for every archetype. On the other hand, as λ increases, households start to borrow at month 0 using borrow_0, and hold the proceeds in savings using save_0. This buffer first appears at $\lambda = 3.50$ for SPA, 1.75 for GWA, and

0.75 for SWA. Because every point is solved as a linear program, each of these thresholds is proven optimal within the model's constraints.

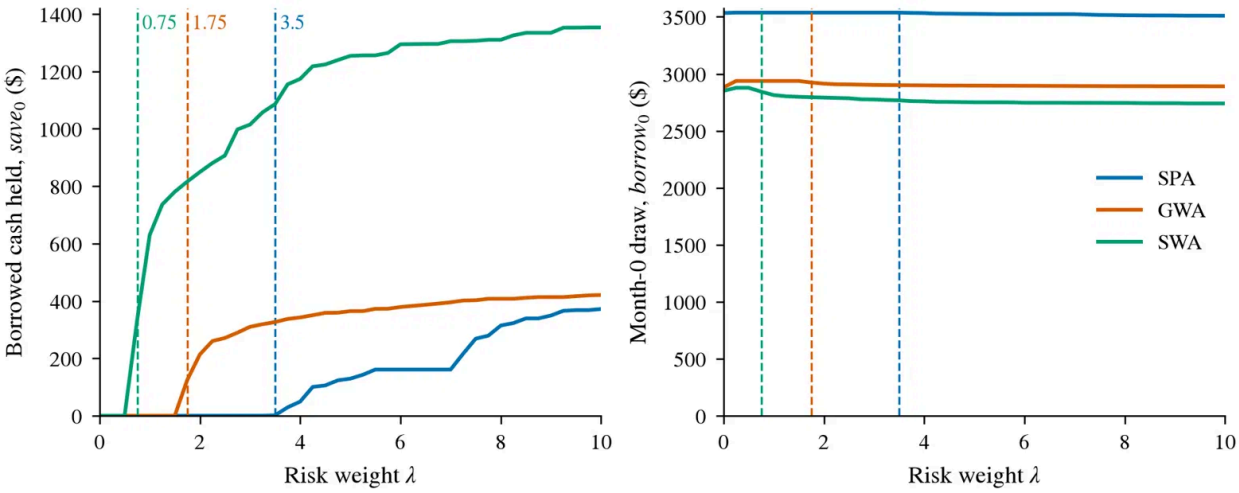


Figure 5 — Single-period save_0 / borrow_0 vs. λ.

The ordering of these households is the result of an archetype's margin over their essential floor. A household that is at higher risk of shortfall is more likely to create a savings buffer to be able to cover the shortfall. For SWA, even a mild concern regarding shortfall justifies paying 21% interest to hold cash. The SWA's floor takes 86% of median income establishing a far larger risk of shortfall. For SPA, that floor sits at 49% of median income; therefore, it has a lower risk of shortfall, meaning that a far stronger concern must exist before the same trade becomes worthwhile. Regardless of the archetype, the buffer is never funded through surplus, instead it is the result of save_0.

Therefore, the mechanism that the multi-period model will test over a 24-month horizon is created by the single-period version's results. The result of this test across debt compounding, savings accumulation, and investment growth is presented in §5.2.

5.2 Optimal Policies

The borrowed cash buffer described in §5.1 almost entirely disappears over the 24-month horizon. In every archetype, most of the surplus goes to investments, between 79% and 89% on average, and the rest to debt repayment. No archetype directs any surplus to savings at $\lambda = 3$. A positive savings share f_S is only present in six of the 360 combinations of archetype, ordering, seed, and λ values; all six occur in a single seed at $\lambda = 20$ or $\lambda = 30$. Savings held from save_0 are also very rare, only occurring at $\lambda = 3$ in one seed, at \$101.

Archetype	f_S	f_I	f_D	borrow_0	save_0
SPA	0.00	0.86 (0.63 - 1.00)	0.14 (0.00 - 0.37)	\$750 (0 - 1,704)	\$20 (0 - 101)
GWA	0.00	0.89 (0.80 - 0.98)	0.11 (0.02 - 0.20)	\$2,108 (0 - 5,353)	\$0
SWA	0.00	0.79 (0.21 - 0.98)	0.21 (0.02 - 0.79)	\$1,679 (0 - 4,801)	\$0

Table 8 — Optimal policy at $\lambda = 3$ (essentials-first; mean and range across five seeds).

In the case of an income drop, households sell investments before borrowing additional funds, meaning that investments act like a buffer in the same way savings would. The assumption of frictionless liquidation of investments, makes this behavior simpler. Additionally, regardless of their variance, investments also earn a market return that is generally larger than r_S in months where they are not needed. In a single month this advantage is insignificant, as investment growth does not have time to develop; however, over a 24-month horizon their accumulation creates the consensus that a separate savings buffer is redundant and slow-growing. Therefore, borrowing at $r_D = 21\%$ to store in save_0 is not worthwhile.

The split between investments and debt, and the size of the initial draw, are less precisely determined. Both of these changes considerably across seeds; for example, SWA's debt share f_D ranges from 0.02 to 0.79. That said, these choices make little difference in the model's outcome: moving the initial draw across its full range only affects the objective by approximately 0.1%, where no difference in pairing is statistically distinguishable from zero. In turn, the model clearly identifies that the household should not hold its surplus in savings; however, when it comes to identifying a set allocation that combines both investments and debt repayment it is quite indifferent.

Two features of the setting favor investments: they are sold without taxes or fees, and the historical return of 12.8% a year is high (§4.5). In §5.10, when the mean return is lowered to 7.0%, the optimal policy's allocation includes a small share of the surplus in savings; however, these policies do not outperform the EV policy, which holds no savings. The finding that surplus should not be held as cash is therefore strongest at historical returns.

5.3 Sensitivity to Risk Aversion

If a household in the model could hold resources to gain protection from bad months, then raising the value of λ would trace a frontier. This is because households that cared more about minimizing tail risk would accept a higher expected cost in exchange for a lower CVaR. Yet, the results do not create a frontier. As λ rises from 0 to 30, the allocation in the optimal policy partially shifts from investments to debt repayment; however, for any archetype, CVaR on the evaluation paths only changes by less than 0.04%, and no such change is statistically distinguishable from zero. On the other hand, the expected

net-worth gain with shifting the allocation decreases significantly for SPA, with $2,688; for GWA and SWA it decreases by $1,029 and $716 respectively, which is not significant. Thus, within the model, more cautious households pay more while gaining no more protection; this is supported by Table 9 which reports the policy and expected costs at $\lambda = 0$, $\lambda = 3$, and $\lambda = 30$.

Archetype	λ	f_I	f_D	$\mathbb{E}[C_s^{fin}]$	**Unmet Need**	$\mathrm{CVaR}_{0.95}$
SPA	0	0.98	0.02	$150,690	$1,818	2,637,014
	3	0.86	0.14	$149,478	$1,818	2,636,950
	30	0.73	0.23	$148,002	$1,817	2,637,979
GWA	0	0.97	0.03	$81,198	$7,273	4,318,627
	3	0.89	0.11	$80,685	$7,281	4,319,436
	30	0.82	0.18	$80,170	$7,284	4,319,620
SWA	0	0.93	0.07	$55,047	$7,693	3,907,601
	3	0.79	0.21	$54,450	$7,701	3,907,232
	30	0.77	0.23	$54,331	$7,704	3,907,598

Table 9 — Policy and costs across λ (essentials-first, mean across seeds).

In Table 9, $\mathbb{E}[C_s^{fin}]$ represents net-worth gain; unmet need represents the expected missed essentials and minimum debt repayments per path over the 24-month horizon. CVaR includes the unmet-need penalty ($M = 100$) discussed in §3.6, meaning it should be read in cost units rather than U.S. dollars. At $\lambda = 30$, one SPA seed also allocates 17% of its surplus to savings.

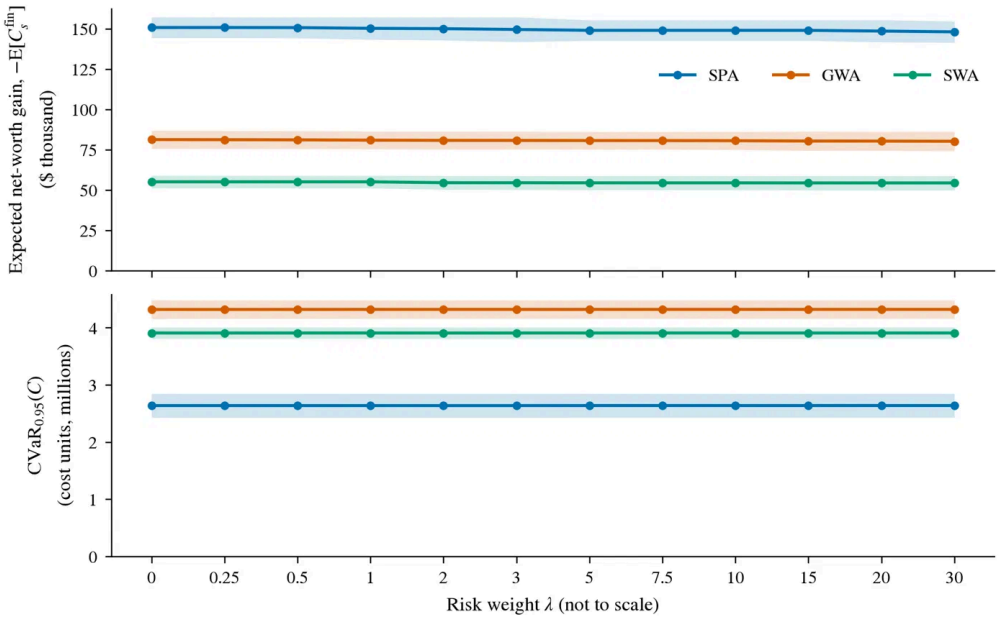


Figure 6 — Net-worth gain and CVaR vs. λ.

The reason this happens is because at every λ value, each of the worst 5% of paths has some unmet need, regardless of archetype or seed. These paths see households spending long stretches, typically a year, in a low income quintile, meaning that no allocation of surplus that rarely even exists can prevent shortfall. In turn, the penalty on that shortfall dominates the cost. Hence, a more cautious policy in stage one cannot reduce shortfall. Table 10 compares the share of paths that have any unmet need across the archetypes. The figures in the table remain nearly unchanged irrespective of the λ value.

Archetype	**Share of Paths with Unmet Need**
SPA	11.9%
GWA	37.4%
SWA	41.0%

Table 10 — Share of paths with unmet need across archetypes.

Therefore, the cost of resilience cannot be measured as a slope along a frontier. Instead, it is measured in §5.5 through a comparison between the optimum and heuristics that build resilience directly by allocating savings.

5.4 Unmet Need and The Tail

For each archetype's cost, unmet need is the largest contributor; in the model, unmet need is predominantly caused by low income rather than inefficient allocation. Table 10 highlights this since at $\lambda = 3$, 11.9% of SPA paths, 37.4% of GWA paths, and 41.0% of SWA paths have at least one month where the household fails to cover its essential spending or minimum debt repayment. Over the 24-month horizon, the expected unmet need is \$1,818, \$7,281, and \$7,701 per path; among the paths that fall short, it averages \$15,250, \$19,454, and \$18,783. Approximately 86% to 89% of it is missed essential spending and the rest is missed minimum debt repayments.

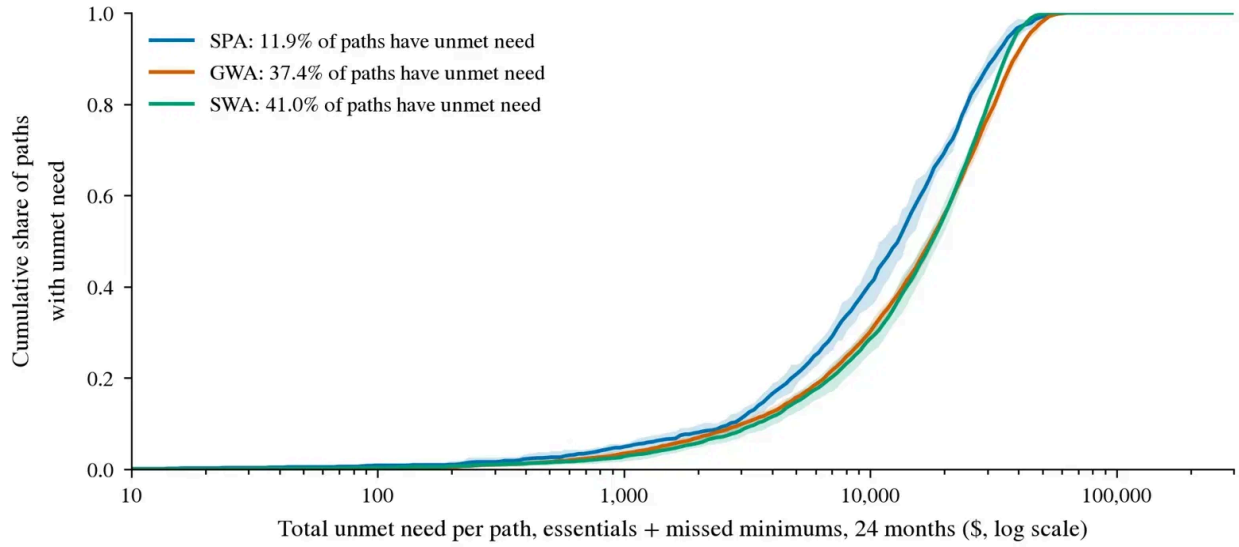


Figure 7 — Unmet-need CDF.

Allocation has a minimal impact on these figures: under every heuristic test the expected unmet need is within \$25 of the optimum's, and the share of paths that fall short is within a quarter of a percentage point. Across the heuristics, the same paths fall short by nearly the same amounts.

Rather than being caused by large shocks in isolated bad months, the shortfalls are mainly the result of sustained low income months. Across the archetypes, months where income is less than 10% of the essential floor are responsible for only: 8% of SPA unmet need, 22% of GWA unmet need, and 6% of SWA unmet need. Instead, most of it is caused by households that spend a year in a low income quintile, where the bottom quintile is below each archetype's annual essential floor; for GWA and SWA, the second lowest quintile is also below said floor. There is little surplus to allocate in such a year, and thus any buffer that exists is exhausted. For any archetype, in any seed, for any value of λ, each of the worst 5% of outcomes has some unmet need.

Three results of this investigation validate the fact that unmet need is mainly caused by income and credit rather than the model's allocation decisions. First, the penalty weight is irrelevant as changing M from 100 to 20 to 500 only moves the unmet need by $2 per path; though, it does change the optimal policy's allocation. Second, increasing the credit limit from three to six months of essential spending improves the objective by up to 30% for SPA, 13% for GWA, and 14% for SWA. The lower objective is almost entirely caused by lower unmet need. Lowering the credit limit to two months worsens the objective by up to 9%; starting with no debt improves it by 4% to 11%. Third, even though credit freezes only affect about 1% of paths where income is near zero, 82% to 91% of frozen months end up with unmet need.

That said, this investigation assumes that a household has no ways of coping outside of withdrawing savings, liquidating investments, or borrowing through their credit. Therefore, these figures should be read as upper bounds for unmet need. In reality, a household can cut spending below the essential floor, borrow from family or friends, or delay bills (§3.10). In turn, a real household facing the same incomes would have fewer and smaller instances of shortfall, though the model effectively demonstrates that unmet need is caused by low income rather than anything else.

5.5 Heuristic Benchmarking

None of the explored heuristics outperform the optimal policy to a statistically significant extent. IF essentially ties it, however, every heuristic that holds savings ends with significantly less wealth. In Table 11, each heuristic is compared to the optimal policy at $\lambda = 3$.

Archetype	Heuristic	$\delta(\pi)$	Net-Worth Gain	Unmet Need
SPA	EFF	0.35%*	-$11,372*	+$10*
	DA	0.56%*	-$11,223*	+$10*
	IF	0.02%	+$1,252	+$0
	FS50	0.39%*	-$11,263*	+$8*
	FS33	0.23%*	-$6,843*	+$3*
GWA	EFF	0.06%*	-$6,764*	+$17
	DA	0.09%*	-$6,627*	+$23*
	IF	-0.03%	+$614*	-$7
	FS50	0.06%*	-$6,680*	+$15
	FS33	0.02%	-$4,096*	+$1

SWA	EFF	0.06%*	-$4,219*	+$13
	DA	0.08%*	-$4,085*	+$25*
	IF	-0.01%	+$793	-$8
	FS50	0.06%*	-$4,134*	+$12
	FS33	0.03%	-$2,392*	+$0

Table 11 — Heuristics vs. optimal policy at $\lambda = 3$.

A statistically significant result within a 95% confidence interval is denoted by *. The results follow essentials-first ordering; differences are the heuristics minus the optimum, per path over 24 months, paired within each seed.

IF is statistically indistinguishable from the optimum in every archetype, and for GWA it accumulates slightly more relative wealth. EFF, DA, and the FS50 split are significantly worse in every archetype; the FS33 split is significantly worse on the objective only for SPA.

The common feature across all losing heuristics is that each one holds savings. In the investigation EFF, FS50, and FS33 are designed to do so directly; however, DA does so indirectly as the heuristic only allocates to savings after clearing its debt balance. Thus, EFF, FS50, and DA forgo almost the same amount of net-worth gain over the 24-month horizon: ~$11,300 for SPA, ~$6,700 for GWA, and ~$4,100 for SWA. FS33 on the other hand, which allocates 33% to investments and only 33% to savings, only forgoes approximately 60% as much as the other three heuristics; this forgone wealth is considered to be the cost of resilience.

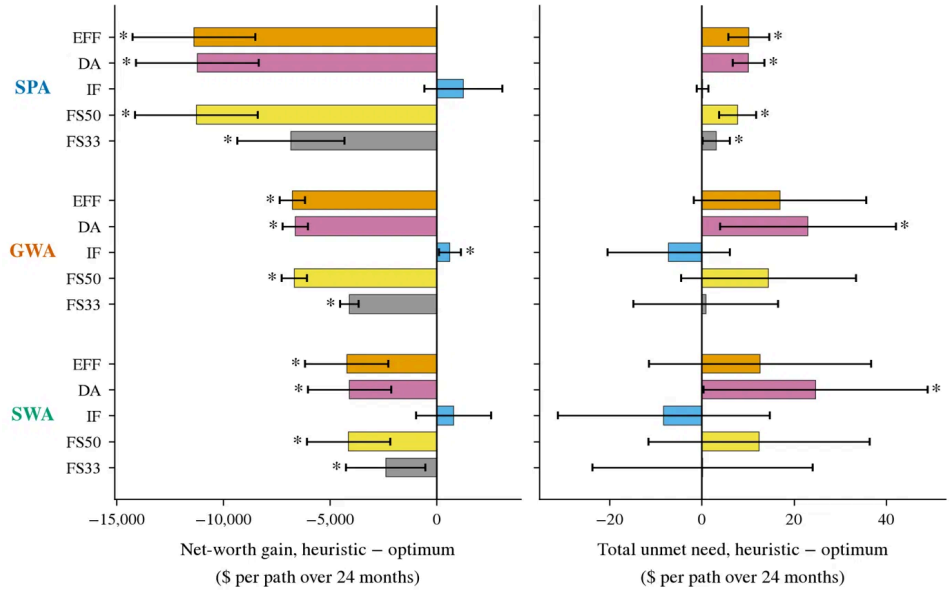


Figure 8 — Difference between heuristics and optimum.

5.6 The Value of Foresight

The bound specification highlights how much lower a household's expected cost could be if it had perfect foresight of the entire income path. If said foresight was given to each household, the objective would improve by 13.2% for SPA, 9.3% for GWA, and 10.4% for SWA. Almost none of this comes from a better allocation: scored with full-foresight recourse, the optimal policy of the main specification is on average within 0.35% of the best full-foresight policy found, and no difference is statistically significant.

Instead, the improvements in the objective are the result of how the full-foresight benchmark handles minimum debt repayments. Though, this does not reflect the behavior of a real household, as no real household has perfect foresight. In the results, unmet essential spending is basically unaffected. Missed minimum debt repayments on the other hand fall by $240 for SPA, $792 for GWA, and $966 for SWA; this fall occurs per path, and brings missed minimum debt repayments to near zero. This result is exemplified by tracing individual paths. In early months, a household with perfect foresight allows essential spending to remain unmet even though credit is available; in turn, this keeps its debt balance below the credit limit. Then, in later months it pays each minimum debt repayment by borrowing the equivalent amount, which leaves the debt balance unchanged. In contrast, the fixed-order household borrows to cover its essential spending, which causes it to reach its credit limit. Thus, it can make no further minimum debt repayments once interest maintains the debt balance above the credit limit. Though, in the case where, in any month, the full-foresight household cannot leave more essentials unmet than the fixed-order household, 96% to 98% of missed minimum debt repayments return.

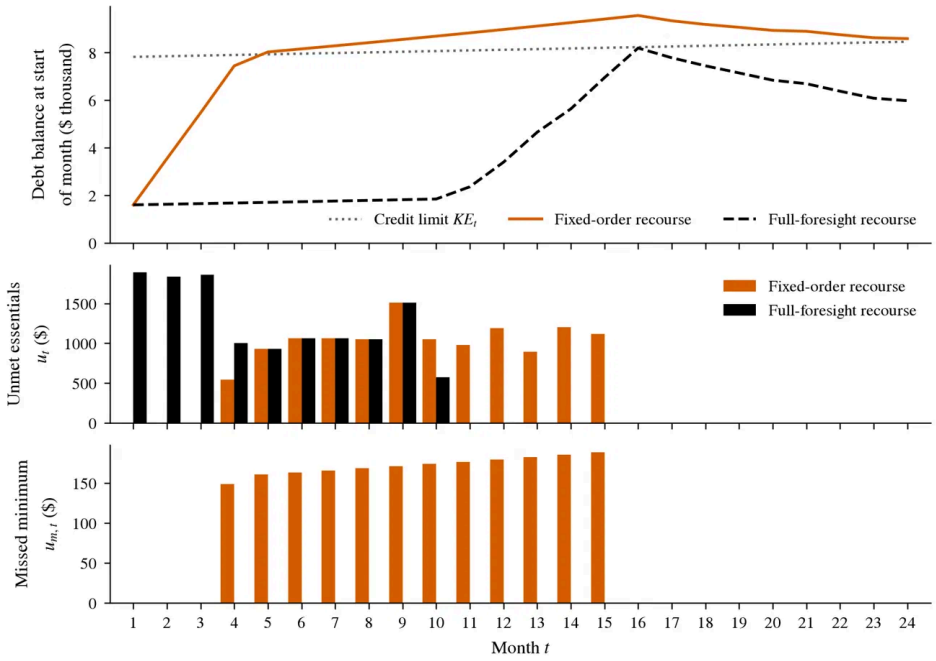


Figure 9 — One simulated GWA path (seed 42).

The reason that the full-foresight household's strategy is appealing to it is the fact that it provides a choice between missed essential spending and missed minimum debt repayments, since both carry the same penalty $M = 100$. Therefore, the only impact foresight actually carries is providing the ability to make that trade rather than missing both. In reality though, no household would trade current essential spending to ensure its minimum debt repayments on a credit card are met in later months. Hence, the foresight gap measures the flexibility of the full-foresight benchmark rather than a value of information that households forgo.

5.7 The Value of the Stochastic Solution and of Perfect Information

The expected-value policy (EV), which plans on a single path of average income, produces a similar result to the optimal policy: it allocates 92% to 96% of its surplus to investments, does not make an initial draw through borrow_0, and does not allocate anything to savings. Though, the EV policy's forecast is very wrong. Because on the average path income never falls below the essential floor, the EV policy predicts no unmet need and an expected net-worth gain $\mathbb{E}[C_s^{fin}]$ of \$158,794 for SPA, \$78,662 for GWA, \$58,450 for SWA across the 24-month horizon. When evaluated on the simulated paths the EV policy results in a similar gain of \$150,717 for SPA, \$81,233 for GWA, and \$55,189 for SWA; though, it leaves 11.9% of SPA paths, 37.4% of GWA paths, and 41.0% of SWA paths with unmet need, resulting in an expected unmet need of \$1,818 for SPA, \$7,273 for GWA, and \$7,693 for SWA per path.

Regardless, VSS is statistically indistinguishable from zero for every archetype, which occurs because the EV policy's objective is statistically indistinguishable from the optimal policy's, thus not affecting the

household's allocation; this also holds when the historical average market return is lowered to 7% for both of these (§5.10). In turn, the only thing that uncertainty should impact is a household's expectations. This happens because a policy that is constructed on the average income picks the right allocation while providing misleading expectations of the risks that come from said allocation.

The EVPI share of the full-foresight objective for each archetype is 2.6% for SPA, 0.9% for GWA, and 1.0% for SWA. EVPI measures only the gain from choosing the allocation separately for each path when the recourse already has full foresight. Because the most important decision, not holding savings, does not depend on which path occurs the EVPI values are small.

5.8 Differences Across Archetypes

The three archetypes all face the same heuristics and arrive at the same conclusions regarding allocations; though, they do differ in terms of exposure. This divergence is less affected by income volatility, as the main thing driving it is how close each archetype's income sits to the essential floor.

Even though SWA income is less volatile than GWA, it is more exposed because its median income is lower than GWA's. Because SWA's essential floor takes up 86% of median income, ordinary income variance is more likely to push SWA below the essential floor; in turn, increasing the share of paths that have unmet need. GWA on the other hand encounters the largest monthly income volatility across the three archetypes and the most months near zero income. Even so, its high-income top quintile results in a higher average income than SWA, which partially offsets expected unmet need. SPA is less prone to unmet need due to its margin, where its essential floor takes up less than half of the median income, meaning that only one in eight paths ever encounters shortfall.

Two results are common to all three archetypes. First, the cost of holding savings is similar in relative terms: EFF gives up 7.6% of SPA's expected net-worth gain, 8.4% of GWA's expected net-worth gain, and 7.7% of SWA's expected net-worth gain. In absolute dollars, this cost is largest for SPA as they have the largest surplus, meaning that misallocation has the largest absolute effect on them. Second, credit access is most important when income is lowest relative to the essential floor. The debt balance is above the credit limit on 11% of SPA paths, 37% of GWA paths, and 40% of SWA paths; for each of these archetypes a larger credit limit reduces unmet need (§5.10).

To summarize, a household's allocation determines how much wealth that household builds, while the percentage of an archetype's essential floor that median income takes up determines how prone it is to unmet need. On the second dimension, the archetypes differ by a factor of three; on the first, the best rule is the same for all of them.

5.9 Stability Across Seeds

Every single result in §5 is reproduced across five independent seeds. For each of these seeds, the policy is selected on 200 optimization paths, and then evaluated on a separate set of 2,000 paths. The intervals reported in §5 reflect the existing variation across the seeds.

The qualitative conclusions made in §7 do not change across seeds. No seed set to $\lambda = 3$ allocates any surplus to savings, IF ties the optimal policy, and the heuristics that hold savings lose wealth. A positive savings share appears only in one seed at $\lambda = 20$ and $\lambda = 30$, and month-0 savings through save_0 only appear in a small number of cells, never above $424.

Though, some smaller details have larger variations across the seeds. For example, across the seeds, there is a large variation in the allocation between investments and debt repayment: SWA's f_D ranges from 0.02 to 0.79. Moreover, borrow_0 ranges from $0 to $5,353. As discussed in §5.2, the objective remains largely unchanged across these choices; therefore, different seeds select different points of a broad optimum rather than entirely different policies. One cell in the robustness checks in §5.10 shows the limit of selecting on 200 paths: at a credit limit of six months, one SWA seed chooses a large initial draw that wins on its optimization paths but loses on the evaluation paths. In turn, this investigation's conclusion in §7 is based on general results that are stable across the five seeds; for example: no savings allocation, IF is near the optimum, etc.

5.10 Robustness

This section tests the main results against any variations in assumed parameters and any variations in parameters that are estimated with uncertainty. Each variant of the said parameters is re-optimized at $\lambda = 3$; this re-optimization follows essentials-first ordering across the five seeds. In Table 12, the change in the objective relative to the main specification is reported.

Variant	**SPA**	**GWA**	**SWA**
Credit limit $K = 2$	+8.9%*	+3.8%*	+4.1%*
Credit limit $K = 4$	-9.5%*	-4.1%*	-4.5%*
Credit limit $K = 6$	-30.0%*	-12.9%*	-14.1%*
$\overline{D_0} = 0$	-11.4%*	-4.5%*	-4.4%*
$\overline{D_0} \cdot 0.7$	-3.4%*	-1.3%*	-1.3%*
$\overline{D_0} \cdot 1.3$	+3.4%*	+1.3%*	+1.3%*

$r_S = 0.5\%$	-0.01%	0.00%	0.00%
Constant Investment Returns	-0.07%	-0.03%*	+0.01%
Income Volatility (low / high)	-0.13% / +0.28%	-0.08%* / +0.07%*	N/A
Freeze Rate ($\pm 20\%$ Window)	0.00%	-0.01%	-0.01%
Excluding 2020-21 Transitions	-0.9%	+0.5%	+0.2%

Table 12 — Change in Z^* relative to the main specification at $\lambda = 3$.

A statistically significant result within a 95% confidence interval is denoted by *. The income volatility is varied in the ranges of relative to σ_{SWA}: 0.25 to 1.00 times for σ_{SPA} and 1.00 to 2.00 times for σ_{GWA}. SWA's freeze rate at its full-year median changes Z by -0.003%.

The results are primarily described by credit access and the initial debt balance $\overline{D_0}$, both of which change outcomes substantially; this happens almost entirely through unmet need. Increasing the credit limit to six months improves the objective for every archetype by 30% for SPA, 12.9% for GWA, and 14.1% for SWA. In turn, this reduces the expected unmet need per path by $823 for SPA, $2,174 for GWA, and $2,356 for SWA. Across the archetypes, starting with $\overline{D_0} = 0$ improves the objective by 4% to 11%. All other parameters only change the objective by less than 0.3%. These include: the savings rate r_S, the volatility of returns, and the freeze rate. Even if the pandemic years are excluded from the model there is no significant change in the objective. Even so, it is important to consider that this check is the least accurate one in this section because missing years means that the regenerated paths cannot be matched precisely with the original ones.

That said, no variant actually changes the main conclusions described in §7. This happens because none of them direct any surplus to savings except for the most extreme penalty weight M, and IF ties the optimum in all 38 combinations of variant and archetype. Moreover, none of the explored heuristics ever beat the optimum by a statistically significant amount: DA is significantly worse in 37 combinations, FS50 in 30, EFF in 28, and FS33 in 21. If penalty variants are excluded the largest gap $\delta(\pi)$ is 0.92% of the objective; this occurs in the DA heuristic for SPA when $K = 6$.

While variants of the penalty weight M lead to the largest changes in the optimal policy, none of them affect the outcomes that matter. When varying M across $M = 20$, $M = 100$, and $M = 500$, the expected unmet need is the same within $2 per path. This supports the notion that unmet need is predominantly impacted by income and credit rather than by how heavily the model punishes it. Furthermore, the only time that varying M impacts f_S is at $M = 500$, where $f_S > 0$ in one seed for two archetypes. Another way that the savings conclusion is weakened is by lowering the mean

investment return to 7%; in doing so, the optimum moves toward debt repayment, and a savings share appears in three of the fifteen seed-archetype cases. That said, none of these policies outperform the EV policy, which holds no savings, on the evaluation paths.

This section is limited by the fact that variants were only optimized at $\lambda = 3$, and a control run found the same optimum as the full search in fourteen of fifteen cases. Additionally, in three seeds (SWA seed 1003 at $K = 6$, SPA seed 1001, and SWA seed 1003 without the pandemic years), the chosen policy of the 200 paths performed poorly on the evaluation paths; in turn, some heuristic comparisons for those variants are considered to be ties.

6 DISCUSSION

6.1 What the Optimal Policy Recommends

The clearest recommendation that the model provides is about what not to do, in the form of: households with uncertain income should not allocate any of their surplus to savings. In the optimal policy, almost none of the surplus is allocated to savings across every archetype, seed, and robustness variant. In contrast, most heuristics that do allocate surplus to savings (EFF, FS50, DA) forgo between $4,100 and $11,400 of their expected net-worth gain over the 24-month horizon. In doing so however, these heuristics do not increase the financial resilience of the households that follow them: their expected unmet need is, if anything, slightly higher.

However, it is important to note that this does not mean a buffer is unnecessary. In the model, the buffer is simply held in the form of investments. This is especially prominent since households in the model liquidate their investments prior to borrowing additional funds; thus, an investment balance V_t covers bad months in the same way that savings would, while earning the historically higher market return. Thus, the precautionary motive identified by Kimball and Carroll does exist within the model; however, unlike in Kimball's and Carroll's models it exists in the form of a liquid investment balance rather than a savings account. The joint holding of savings and credit card debt documented by Gross and Souleles has been explained by cash-only expenses (Telyukova), precautionary borrowing (Druedahl and Jørgensen), and volatile credit limits (Fulford). These mechanisms are largely missing from the model, which is one of the reasons that savings are almost never optimal here.

Other than that, the model remains largely indifferent to allocations. For example, the outcome is barely changed if the split between investing and repaying debt changes; investing the entire surplus is statistically indistinguishable from the optimum. Additionally, greater caution through the form of higher λ values does not provide the household with any additional benefit. The reason for this is that the shortfalls that households fear occur as a result of income falling below the essential floor, and the only factors that actually reduce the number of shortfalls are a larger credit line or less starting debt, both of which lie outside of the model's allocation decision.

It is crucial to note that these recommendations depend on assumptions that favor investing; in the model, the liquidation of investments occurs without taxes or fees, returns are independent of income, and a high historical market return of 12.8% all make investing more appealing. At a 7% mean market return, the optimum shifts towards repaying the 21% interest rate credit card debt. Furthermore, a household that is facing friction in liquidating investments through the form of taxes or sale delays, or a household in a recession that reduces both its income and the value of its investment portfolio simultaneously, would have more of a reason to hold a savings buffer. Hence, the recommendations made by the model should be read as a ranking that assumes favorable conditions. Under this ranking, savings are the costliest way of holding a buffer, and the rankings of investments and debt repayment depend on the mean market return.

6.2 Implications for Financial Education

Within the financial literacy industry, a common rule of thumb given to households is to build an emergency fund that covers three to six months of essentials (Scott). That said, within the model that recommendation is almost never optimal. It takes SWA approximately nineteen months at its monthly median income to save up three months of essentials. In the 24 months modelled in the multi-period version, EFF causes households to forgo between $4,219 and $11,372 of their expected net-worth gain, without decreasing their unmet need. The instinct to hold a buffer is correct, but doing so in the form of savings is not. Within this investigation's limitations (§6.3), a liquid investment balance V_t accomplishes the same goal at a cheaper cost. Therefore, financial education should cover the option of holding a buffer in the form of an investment portfolio. In doing so, it should also explain when doing so can fail for many households; such instances include when liquidation is taxed or delayed or when an economic downturn hits income and investments together.

Additionally, heuristic selection is far less impactful than the financial literacy industry suggests. Every single one of the tested heuristics is within approximately 1% of the optimum's objective; none of the heuristics change the share of paths that experience shortfall by more than a quarter of a percentage point. When considering shortfall, the two factors that matter most are the percentage of income necessary to cover essential spending and credit access. As demonstrated in §5.6, households that borrow to be able to afford essential spending quickly reach their credit limit; after that happens, interest keeps D_t above the credit limit and every minimum debt repayment is missed. It is thus more important to maintain low debt in order to preserve the capacity of the credit line than to decide which heuristic to use.

While using the EV policy may seem logical to some since it does provide the correct allocation, doing so fails to accurately predict shortfall for households whose actual chance of unmet need over two years is 37% to 41%. Because many borrowers do not have the ability to evaluate the advice they receive (Lusardi and Tufano), showing them the range of likely outcomes would give a more honest picture of the decision they face.

6.3 Limitations

Limitation	Likely Effect
Model Structure	
Shares are fixed for 24 months.	Understates what an adaptive household could achieve.
Shortfalls are covered in a fixed order.	Overstates unmet need relative to smarter month-to-month choices.
Full-foresight benchmark trades unmet essentials for minimum payments.	Overstates the value of foresight (§5.6)
Minimum debt repayment cannot be paid from current income.	Raises missed minimum debt repayment, especially early in the 24-month horizon.
The 24-month horizon.	Long-run compounding of investments and debt is not captured.
Household Behavior	
No spending cuts, family help or other credit.	Unmet need is an upper bound.
No discretionary spending.	Overstates net-worth accumulation.
Single-person, single-earner household.	Ignores pooled incomes and dependents; the net effect is ambiguous.
One policy and one set of starting conditions per archetype.	Results describe a typical household, not the spread within each archetype.
Financial Assumptions	
Investments sold without taxes, fees, or delays.	Favor investing over saving.
High historical return (12.8% per year).	Favor investing; at 7% the optimum shifts toward debt.
Returns drawn independently each month.	Crashes do not cluster, favoring investing.
Returns drawn independently of income.	Crashes do not coincide with income losses, favoring investing.

Fixed interest rates.	Minor over 24 months.
A single credit card debt.	Debt avalanche and debt snowball coincide; other types of debt are ignored.
No over-limit fees or penalties.	Understates the cost of carrying debt above the limit.
Calibration	
Credit limit K and the 2% minimum debt repayment are assumed.	Outcomes are sensitive to K (§5.10).
Penalty weight M is assumed.	Changes the policy but not unmet need (§5.10).
Volatility of GWA and SPA set relative to SWA.	Assumed, not estimated; effect under 0.3% (§5.10).
Freezes assigned to the lowest-income months, at quintile-matched rates.	Exogenous, not triggered by borrowing or missed debt repayments; little effect (§5.10).
Incomes are nominal while the floor rises with inflation.	Real incomes fall about 8%, adding 1.5-1.9 points to below-floor months.
SCF debt is matched on household income.	Includes multi-earner households; approximate match.
GWA covers all self-employment; median income overstated by about 9%.	Not specific to the platform of work; GWA is slightly better off than its reality.
Pandemic years are included in the transitions.	Excluding them changes no conclusion, but the check is imprecise (§5.10).
Solution Method	
Policies chosen on 200 paths.	Some seeds overfit; conclusions rest on results stable across the seeds (§5.9).
Grid search rather than a proven optimum.	Optimal only to within the grid resolution.
13 of 60 bound searches stopped early; HiGHS settings sometimes disagree.	The full-foresight benchmark may be slightly loose.

Table 13 — Limitations and their likely effect on the results.

7 CONCLUSION

This investigation asked how a household with uncertain income should divide its monthly surplus between savings, investments, and debt repayment. In doing so, it also inquired about what it costs to prioritize financial resilience. In order to answer the research question, this paper formulates the decision as a two-stage stochastic program with a CVaR term; the data within it is calibrated to federal microdata for three archetypes. After the formulation and calibration are completed, the optimal policy is benchmarked against the heuristics that are most commonly recommended by the personal finance industry; this is done under the same income paths and the same fixed-order recourse.

This investigation's central finding is that a buffer held in the form of savings reduces the expected net-worth gain of a household, while purchasing no additional protection from bad income and market months. Within the optimal policy, the surplus was almost never allocated to savings, and investing the entire surplus is statistically indistinguishable from it in the objective. Heuristics that hold savings like EFF, FS50, and DA, forgo between $4,100 and $11,400 of expected net-worth gain over the 24-month horizon. Following these heuristics also results in slightly higher unmet need, by $8 to $25 per path, meaning that saving purchases no additional protection. The cost of resilience in this investigation is represented by the forgone wealth. However, no frontier between expected cost and tail risk can be traced because greater caution, through larger λ values, increases the expected cost without reducing CVaR. The reason for this is that within the tail's paths, income falls below the essential floor, meaning that no allocation actually offsets tail risk.

Even though these conclusions hold across five distinct seeds and robustness variations across different parameters, they do rely on certain assumptions that could be removed in future work. First, returns could be drawn in historical sequences and linked with income, which would cause economic downturns to coincide with job loss or wage cuts. Second, the liquidation process for investments could be made less frictionless through taxes, fees, and potential sale delays. Third, causing freezes and penalties to follow from the household's past borrowing and missed debt repayments could make them more accurate. Fourth, households could be allowed to cut spending, have discretionary consumption, informal borrowing (from friends or family), and second earners. Fifth, the problem could be presented as a multi-stage formulation that allows households to change their shares depending on their current balances. Sixth, the penalties on missed essential spending and minimum debt repayments could be distinct, reflecting the gravity of each and removing the artifact in the full-foresight benchmark.

Within the model's limitations, this investigation finds a clear answer to the research question. The cost of resilience is very real; however, for it to be exemplified clearly a household needs a margin between its income and essential floor.

REFERENCES

Artzner, Philippe, et al. "Coherent Measures of Risk." *Mathematical Finance*, vol. 9, no. 3, 1999, https://doi.org/10.1111/1467-9965.00068.

Board of Governors of the Federal Reserve System. "Commercial Bank Interest Rate on Credit Card Plans, All Accounts (TERMCBCCALLNS)." *FRED, Federal Reserve Bank of St. Louis*, fred.stlouisfed.org/series/TERMCBCCALLNS. Accessed 7 July 2026.

---. *Survey of Consumer Finances, 2022*. Federal Reserve, 2023, www.federalreserve.gov/econres/scfindex.htm. Accessed July 2026.

Carroll, Christopher D. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics*, vol. 112, no. 1, 1997, https://doi.org/10.1162/003355397555109.

Carroll, Christopher D., and Lawrence H. Summers. "Consumption Growth Parallels Income Growth: Some New Evidence." *National Saving and Economic Performance*, edited by B. Douglas Bernheim and John B. Shoven, U of Chicago P, 1991, pp. 305–48.

Consigli, Giorgio, et al. "Long-Term Individual Financial Planning under Stochastic Dominance Constraints." *Annals of Operations Research*, vol. 292, no. 2, 2020, https://doi.org/10.1007/s10479-019-03253-8.

Deaton, Angus. "Saving and Liquidity Constraints." *Econometrica*, vol. 59, no. 5, 1991, https://doi.org/10.2307/2938366.

Druedahl, Jeppe, and Casper Nordal Jørgensen. "Precautionary Borrowing and the Credit Card Debt Puzzle." *Quantitative Economics*, vol. 9, no. 2, 2018, https://doi.org/10.3982/QE621.

Federal Reserve Bank of New York. "30-Day Average SOFR (SOFR30DAYAVG)." *FRED, Federal Reserve Bank of St. Louis*, fred.stlouisfed.org/series/SOFR30DAYAVG. Accessed 7 July 2026.

Flood, Sarah, et al. *IPUMS CPS: Version 13.0*. IPUMS, 2025, https://doi.org/10.18128/D030.V13.0.

French, Kenneth R. "Fama/French 3 Factors." *Kenneth R. French Data Library*, Tuck School of Business, Dartmouth College, mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. Accessed 7 July 2026.

Fulford, Scott L. "How Important Is Variability in Consumer Credit Limits?" *Journal of Monetary Economics*, vol. 72, 2015, https://doi.org/10.1016/j.jmoneco.2015.01.002.

Gal, David, and Blakeley B. McShane. "Can Small Victories Help Win the War? Evidence from Consumer Debt Management." *Journal of Marketing Research*, vol. 49, no. 4, 2012, https://doi.org/10.1509/jmr.11.0272.

Ganong, Peter, et al. "Earnings Instability." Working Paper 34227, National Bureau of Economic Research, 2025, rev. July 2026, www.nber.org/papers/w34227.

Gross, David B., and Nicholas S. Souleles. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics*, vol. 117, no. 1, 2002, https://doi.org/10.1162/003355302753399472.

Huangfu, Qi, and J. A. Julian Hall. "Parallelizing the Dual Revised Simplex Method." *Mathematical Programming Computation*, vol. 10, no. 1, 2018, https://doi.org/10.1007/s12532-017-0130-5.

Kimball, Miles S. "Precautionary Saving in the Small and in the Large." *Econometrica*, vol. 58, no. 1, 1990, https://doi.org/10.2307/2938334.

Konicz, Agnieszka Karolina, et al. "A Combined Stochastic Programming and Optimal Control Approach to Personal Finance and Pensions." *OR Spectrum*, vol. 37, no. 3, 2015, https://doi.org/10.1007/s00291-014-0375-6.

Lusardi, Annamaria, and Peter Tufano. "Debt Literacy, Financial Experiences, and Overindebtedness." *Journal of Pension Economics and Finance*, vol. 14, no. 4, 2015, https://doi.org/10.1017/S1474747215000232.

Mitchell, Stuart, et al. *PuLP: A Linear Programming Toolkit for Python*. U of Auckland, 2011.

Morduch, Jonathan, and Rachel Schneider. *The Financial Diaries: How American Families Cope in a World of Uncertainty*. Princeton UP, 2017.

Rockafellar, R. Tyrrell, and Stanislav Uryasev. "Optimization of Conditional Value-at-Risk." *Journal of Risk*, vol. 2, no. 3, 2000, https://doi.org/10.21314/JOR.2000.038.

Scott, Cindy. "An Emergency Fund Is Key to Your Financial Strategy." *Charles Schwab*, 11 Nov. 2025, www.schwab.com/learn/story/why-an-emergency-fund-is-essential.

Telyukova, Irina A. "Household Need for Liquidity and the Credit Card Debt Puzzle." *Review of Economic Studies*, vol. 80, no. 3, 2013, https://doi.org/10.1093/restud/rdt001.

U.S. Bureau of Labor Statistics. "Consumer Price Index for All Urban Consumers: All Items in U.S. City Average (CPIAUCSL)." *FRED, Federal Reserve Bank of St. Louis*, fred.stlouisfed.org/series/CPIAUCSL. Accessed 7 July 2026.

---. "Table 3404. One Person Consumer Units by Income before Taxes, Consumer Expenditure Surveys, 2023–2024." *Consumer Expenditure Surveys*, www.bls.gov/cex/tables.htm. Accessed 1 Aug. 2026.