

Counting Hours, Counting Losses: The Toll of Unpredictable Work Schedules on Financial Security

Anonymous authors

Paper under double-blind review

Abstract

Financial instability is a pressing concern in the United States, with drivers that include growing employment disparities and insufficient wages. While research typically focuses on financial aspects such as income inequality in precarious work environments, there is a tendency to overlook the time-related aspect of unstable work schedules. The inability to rely on a consistent work schedule not only leads to burnout and conflicts between work and family life but also results in financial shocks that directly impact workers' income and assets. Unforeseen fluctuations in earnings pose challenges in financial planning, affecting decisions regarding savings and spending, and ultimately undermining individuals' long-term financial stability and well-being. This issue is particularly evident in sectors where workers experience frequently changing schedules without sufficient notice. The lack of advance notice disproportionately affects vulnerable groups, including those in the food service and retail sectors, part-time and hourly workers, individuals with lower incomes and education levels, and specific racial groups. These groups are already more financially vulnerable, and the unpredictable nature of their work schedules exacerbates their financial fragility.

Our objective in this study is to understand how unforeseen fluctuations in earnings exacerbate financial fragility by investigating the extent to which individuals' financial management depends on their ability to anticipate and plan for future events. To address this question, we develop an online learning approach in which individuals adapt their consumption strategies over time in response to financial uncertainty and evolving information. This approach forms the basis of our simulation framework, which models how workers manage consumption in the face of variable work schedules and the imperative to avoid financial ruin.

With this framework, we demonstrate both theoretically and empirically how a worker's capacity to anticipate schedule changes enhances their long-term utility. Conversely, the inability to predict future events can worsen workers' financial instability. Moreover, our framework enables us to explore interventions aimed at mitigating the problem of schedule uncertainty and evaluate their effectiveness.

1 Introduction

Financial insecurity has become a defining feature of modern life in America (Weller, 2018). A convergence of socioeconomic factors, escalating inequality, and precarious employment situations (Kalleberg, 2009; 2011) have rendered families more susceptible to financial shocks with irreversible long-term consequences.

Unpredictable work schedules which are widespread in many sectors like the food-service and retail sectors, play a pivotal role in this insecurity. In addition to low wages and limited benefits, economically disadvantaged workers often contend with unpredictable work schedules. Previous research reveals that 80% of food industry workers have minimal to no influence on their schedules, and 69% are obligated by the system to maintain schedules that are "open and available" for work at any given time (Loggins, 2020; Schneider & Harknett, 2019b). 8.4% of workers and consultants/contractors aged 18-65 reported significant fluctuations in their income monthly. 51% attributed their instability to irregular work schedules (McCrate, 2018).

With a just-in-time work schedule, managing one’s finances becomes increasingly difficult. Workers earn income through their employment and use it to fulfill daily necessities like food, shelter, transportation, clothes, recreation, and so on. In this context, the income earned by workers is allocated to different purposes, leading to corresponding gains in utility based on how much they save and consume optimally. Given that work schedule instability has a direct impact on individuals’ employment and income, an important line of inquiry arises: What are the adverse effects of this scheduling discrepancy on these protected groups? Additionally, how can these potential adverse impacts be effectively measured?

To answer these questions, we need to study the resultant consumption and saving behavior of workers in the face of unforeseen financial shocks from unstable shifts. Thus the primary goal of this paper is to thoroughly investigate the dynamics/effects of work instability on the earned utility of workers. To do this, we need an agent-based behavioral model capable of simultaneously simulating several interconnected phenomena, i.e., the model should encompass how individual consumption responds to unexpected financial shocks and the behavior of individuals as they seek to maximize utility with different levels of information on future events.

However, most economic models of consumption and savings assume that individuals possess the ability to fully look into the future when making consumption decisions (Deaton et al., 1992; Deaton, 1989). On the contrary, in real life, information is limited, and our ability to act on it is also limited. Thus, we need to understand how individuals’ financial decision-making changes if they only have a limited window of information into how their financial well-being might change. This is particularly relevant in settings where workers are increasingly only given limited visibility into income/asset-affecting decisions like how many hours they will work or at what rate they will earn.

In this paper, we seek to model this problem using the language of online learning, i.e., an adaptive update of the workers’ consumption policies (an online learning framework) where the policy is recalculated at each step as more information on work schedules becomes available. We investigate how far utility maximizing strategies depend on the degree of “foresight”. The modeling approach is advantageous for several reasons. First, offline models generate a singular, fixed policy incapable of adapting to new financial information/shocks. Secondly, the simultaneous occurrence of financial data availability and optimization better aligns with real-world behavior. Third, the introduction of future lookahead allows for a nuanced exploration of its utility under different conditions and parameters (i.e., various levels of foresight). Lastly, the online framework with future information provides a robust simulation environment for studying the consequences of work schedule instability and bias in future information availability.

It also opens avenues for empirically exploring mitigations. This is in line with the current ongoing efforts in adopting various policies and regulations to make altering work schedules more equitable, e.g., the San Francisco Board of Supervisors’ mandates (based on Retail Workers Bill of Rights) that ensure more advance notice for hourly workers in retail stores when establishing or altering work schedules (Golden, 2015).

Contributions. Overall, the main contributions of this work are:

- **Online Learning Algorithm with Lookahead:** A novel online algorithm is proposed, capable of handling varying levels of lookahead. The algorithm allows individuals to utilize the deterministic information from the lookahead to modify their consumption decisions in real-time, responding to and adapting to financial shocks as they occur. This algorithm becomes a valuable tool for studying the impact of different degrees of foresight on decision-making.
- **Formal Analysis of the Effects of Lookahead:** Workers who have a lookahead benefit from an advantage that increases proportionally with the magnitude of their lookahead compared to workers with no lookahead. It is also important to note that this gap is tight under appropriate assumptions.
- **Empirical Analysis of Future Information Provision:** Individuals (workers) are provided with future information on certain upcoming events. This aspect enables an empirical exploration of the extent to which future information aids in financial management, contributing to increased utility.
- **Temporal Equity in Workplace Schedules:** We explore temporal equity, particularly in the context of the implications of lack of advance notice (future lookahead) on work timetables that affect the disadvantaged subpopulations more acutely.

- **Mitigation Strategies:** Various mitigation strategies with reference to fair workplace laws and acts are explored to understand how the adverse effects of just-in-time work schedules, particularly in terms of workers’ utility, can be alleviated. This investigation aims to identify effective strategies for improving the overall well-being of workers facing schedule instability.

2 Related Work

Work Schedule Instability. The current focus of research primarily lies in the field of sociology, specifically examining irregular work scheduling and its various repercussions. Unstable schedules cause income volatility (Hannagan & Morduch, 2015; Morduch & Siwicky, 2017; Farrell & Greig, 2016; Reserve, 2016; Schneider & Harknett, 2017) and income volatility results in financial and life hardship (Bania & Leete, 2006; Reserve, 2016; Leete & Bania, 2010; McCarthy et al., 2018; Lambert et al., 2019). This encompasses issues such as burnout from precarious work schedules (Schneider & Harknett, 2019a; Hawkinson et al., 2023) and work-family conflicts (Golden, 2015; Henly & Lambert, 2014), particularly affecting parents with unpredictable or just-in-time schedules. The impact also extends to areas like anxiety and child behavioral problems linked to parental work instability (Schneider & Harknett, 2019b). Additionally, the field of Human-Computer Interaction (HCI) has also strived to study similar repercussions with a participatory outlook (Uhde et al., 2020; Wood, 2021; Lee et al., 2021). They investigate the necessity of worker participation in deciding schedules to ensure fairness (Uhde et al., 2020) and how to use elicitation methods to figure out worker preference models to assist in schedule management (Lee et al., 2021).

Further, statistical data shows a pronounced unfairness in advance notices for altering work timetables for underprivileged groups. Hourly workers, individuals with lower educational attainment, women of color, and specific service sectors are disproportionately affected (Schneider & Harknett, 2019b; McCrate, 2018) with managerial discretion (Wood, 2018; Lambert, 2008). These work schedules make it difficult to plan for the future (Shah et al., 2015) and difficulty in planning would result in disproportionate financial poverty and hardship (Gennetian & Shafir, 2015). Along the same lines, there are some reports pointing to scheduling software and planning algorithms as a factor behind more unpredictable scheduling, particularly for low-wage workers in the service industry (Kantor, 2014; Lancaster, 2015; Griesbach et al., 2019; Zhang et al., 2022). For example, a New York Times article pointed out how some employees with algorithmic schedules rarely learned their timetables more than three days before the start point of a workweek (Kantor, 2014) or how pay reduction is correlated with sudden schedule changes and sales figures (Loggins, 2020).

Simulation. Simulation has been used in many social settings, including fairness in lending (Liu et al., 2018), resource allocation (Ensign et al., 2018; Elzayn et al., 2019), college admissions (Hu et al., 2019; Kannan et al., 2019), financial analysis (Cristelli et al., 2011), technology adoption (Klügl & Kyvik Nordås, 2023), studying supply chain shortages (Yongsatianchot et al., 2023), and simulations of global crisis like the pandemic to minimize the spread of virus (Abe et al., 2022). To study the effects of advance notice on individuals, we follow in the path of research by D’Amour et al. (2020), Zheng et al. (2020), and Nokhiz et al. (2021), which use simulation as a mechanism to study behaviors of agents in systems.

Consumption Models. Consumption models fall within the broader category of discounted utility (DU) models that collectively contribute to our understanding of how individuals navigate decisions related to consumption and saving over time (Deaton et al., 1992; Deaton, 1989). These models involve individuals engaging in discounted consumption utility maximization, where they generally prefer immediate rewards over future rewards of the same size (Chabris et al., 2010). The decision-making process involves choices regarding when to consume or save (Samuelson, 1937). Several widely used intertemporal consumption models include the permanent income hypothesis (PIH), the life-cycle model (Deaton et al., 1992; Friedman, 2018; Parker, 2010), and the neoclassical consumption model (Bütler, 2001). The income fluctuation problem (IFP) is also a related consumption model with an infinite time horizon optimization with income uncertainty and an upper bound on consumption (the upper bound hinders any consumption more than the amount of assets individuals own, i.e., one cannot consume more than what they possess) (Ma et al., 2020; Sargent & Stachurski, 2014; Deaton, 1989; Den Haan, 2010; Kuhn, 2013; Rabault, 2002).

Online Learning in Investment. In investment, where high volatility and real-time data prevail, an online model is needed to allocate investments across assets and maximize cumulative wealth through sequential

optimizations. This application of online learning, known as online portfolio selection (Dochow, 2016; Xi et al., 2023; Li & Hoi, 2014; 2018), is an algorithmic trading strategy that predicts future prices of risky assets using historical data. Online learning algorithms then optimize the portfolio with loss functions tailored to financial objectives, aiming for maximum wealth.

While investment models (Bayraktar & Young, 2012; Grandits, 2015) have valuable components like modeling of uncertainty, they can only imperfectly model consumption. They focus on portfolio allocation between risky and risk-free assets rather than consumption decisions. Although individuals can adjust allocations to avoid financial ruin, this has no direct equivalent in everyday saving and spending choices. Moreover, investment models do not impose consumption constraints.

Reinforcement Learning and Lookahead. The concept of lookahead has garnered significant attention in the recent landscape of reinforcement learning (RL). Generally, in RL models, recent works look into the idea of H -step lookahead. In H -step lookahead (Sikchi et al., 2022), the learner has a learned dynamic model and that is used to calculate the action sequence to a horizon of size H to find the optimal policy that maximizes the cumulative result. There are also approaches where the learner incorporates a greedy real-time dynamic programming algorithm, replacing the greedy step with an H -step lookahead policy (Efroni et al., 2020). The works of (Moerland et al., 2020; Sikchi et al., 2022; Dalal et al., 2021; Efroni et al., 2018) study the properties of H -step lookahead where the horizon is of a fixed size. Compared to the H -step lookahead in RL which uses the expected utility summed over H future steps, our approach deterministically reveals all information across that fixed H -step horizon. Furthermore, in RL models explored in these works, there is a reliance on explicitly defined states and actions, involving the learning of model dynamics on the state-action space, utilizing lookahead. This learned information is then used in sampling the next state. In contrast, our models operate with no *explicitly* defined states and actions (even though there is an implicit state space due to assets).

The recent works of Merlis et al. (2024); Merlis (2024) study competitive ratios and regrets of rewards under a notion of partially deterministic lookahead. In these works, the authors explore RL problems where the agents have access to information on the rewards of their actions or the state transitions before an action is made at the current time point, and they explore how this lookahead information helps the agents improve their final utility. The assumption of known rewards and transitions can be seen as an intermediate notion of lookahead lying between the standard RL lookahead, which relies on learned estimates, and our formulation, where future information is assumed to be known exactly.

3 Determining Consumption and Savings: An Adaptive Online Learning Algorithm

A crucial component of a framework designed to investigate the significance of advance notice (future lookahead) involves creating a model that captures consumption under lookahead while accommodating uncertainty – specifically, how individuals, facing financial uncertainty, make decisions regarding the amount to consume and save. Although various models attempt to represent consumption under uncertainty, they all fundamentally rely on the concept of discounted utility, which is the most common model in economics for understanding the interplay between consumption and savings. In a discounted utility model, the agent, at each time step, consumes a certain amount c and receives utility $u(c)$ from some concave function u . The objective is to devise a policy to determine a consumption value c in a way that maximizes the total discounted utility. After this brief introduction, we will formally articulate the specific cases we study, assumptions, and models in the following section (§3.1).

3.1 Our Models

In this section, we introduce the main models that we study in this work. We consider both deterministic and stochastic models that use a discount factor β to compute the utility. We assume that time is discretized into integer steps, and let T be the effective overall job timeframe, i.e., the algorithm’s time horizon. Let a_t, c_t, y_t be the assets, consumption, and income at time t , respectively. We also let R_t denote the gain from assets (independent of the size of the assets), i.e., the appreciation/depreciation rate of assets. Let $u(\cdot)$ (in

our context, we employ $u(c) = \sqrt{c}$ which is in the class of isoelastic utility functions that are commonly used in macroeconomics) be the utility. Let \mathcal{D}_Y be the income distribution and \mathcal{D}_R be the returns distribution.

In all the models, the goal of the algorithm is to maximize the total utility. More formally, it is to solve the following optimization problem:

$$\begin{aligned} \max_{c_1, \dots, c_T} \quad & \sum_{t=1}^T \beta^{(t-1)} u(c_t) \\ \text{subject to:} \quad & a_{t+1} = R_t(a_t - c_t) + y_t \text{ and } 0 \leq c_t \leq a_t \end{aligned} \quad (1)$$

The constraint $0 \leq c_t \leq a_t$ ensures that the worker consumes from the assets available to them. This constraint ensures the worker could consume without going to ruin. The model equation $a_{t+1} = R_t(a_t - c_t) + y_t$ shows how assets evolve over time given the consumption and income.

The offline or deterministic model. In the offline model, the income and return values y_t, R_t are known a priori to the algorithm (as, of course, is the starting asset value, a_1). This is the most common model considered in traditional economic theory, and the solution can be found using dynamic programming. The “states” in the dynamic programming correspond to the time step of interest and the total assets remaining.

The stochastic model. In the stochastic model, the income and return values are stochastic and they come from known distribution. In this case, an algorithm can optimize the modified objective:

$$\max_{c_1, \dots, c_T} \sum_{t=1}^T \beta^{(t-1)} \mathbb{E}(u(c_t)) \quad (2)$$

The parameters and the constraints follow the same definitions as in the deterministic case, with the caveat being that y_t, R_t are updated using the “realized” values (not their expectations). Note that this is already an online algorithm: at every t , the algorithm computes the consumption value using the expected y, R for the future, but as the y_t, R_t values get updated, the algorithm changes its behavior.

Finally, to explore the effects of future information (lookahead), we consider a combination of the two models.

3.2 Online Consumption Algorithm with Lookahead

We finally consider a model that has “limited determinism” controlled by the extent of lookahead, and the rest of the process is stochastic. Formally, at any time t , the algorithm knows the exact values of y_{t+i} and R_{t+i} , for $i = 1, 2, \dots, \tau$, for some lookahead parameter τ .¹ Additionally, it knows the distribution of the parameters for later time steps. Once again, in contrast to offline economic models, our online model finds an adaptive consumption policy that changes over time given the lookahead as more information arrives.

The algorithm itself is a straightforward combination of the deterministic and stochastic cases. At each step, the algorithm computes:

$$\max_{c_1, \dots, c_T} \sum_{t=0}^{\tau} \beta^t u(c_t) + \mathbb{E} \left(\sum_{t=\tau+1}^T \beta^t u(c_t) \right) \quad (3)$$

where \mathbb{E} is the expectation over the stochasticity in income and returns beyond the lookahead horizon.

Note that the algorithm incorporates new data (as well as lookahead information) as it arrives, which is why it is an online learning algorithm. For simplicity, we assume that a lookahead of τ implies the agent is aware of the exact return and income values for the next τ time steps, and for the time steps beyond τ , distributions of these parameters are known.

¹It is also interesting to consider further hybrid models; e.g., the algorithm has a lookahead over y_t , but not over R_t . This is also realistic in practice since return rates are governed by the market while income lookahead can be controlled by the employer. We do not consider such models in this work.

Algorithm 1 Online consumption with lookahead

- 1: Let τ be the lookahead, a_0 be the initial assets
- 2: Solve the following optimization problem and save the maximum utility for each a, t in a table V where $V[a, t]$ gives the utility of a, t pair,

$$\max_{c_1, \dots, c_T} \mathbb{E} \left(\sum_{t=0}^T \beta^t u(c_t) \right)$$

$$x_{t+1} = R_t(a_t - c_t) + y_t \text{ for all } t \text{ } (y_t \in \mathcal{D}_Y, R_t \in \mathcal{D}_R \text{ for } t \in \{0, \dots, T\})$$

- 3: **for** $r = 0$ to T **do**
- 4: Solve the following optimization problem to get c_r given a_r using dynamic programming,

$$\max_{c_1, \dots, c_T} \sum_{t=r}^{r+\tau} \beta^t u(c_t) + V(a_{r+1+r}, \tau + 1 + r)$$

$$a_{t+1} = R_t(a_t - c_t) + y_t \text{ for all } t \text{ } (y_t, R_t \text{ exactly known for all } t \in \{r, r+1, \dots, r+\tau\})$$

- 5: Set $a_{r+1} = R_r(a_r - c_r) + y_r$
- 6: **return** The sequence of c_t s at each time t

For this model, we define Algorithm 1. This algorithm solves a stochastic dynamic programming (DP) at the start (in Line 2) where it populates the table V (of size $* \times T$ where $*$ is a placeholder for the size of the assets, with dynamic resizing as the bounds for the assets increase) with the maximum utility values for each asset and time point pair in the possible paths of the stochastic system. The algorithm uses backward induction over a discretized space using the known distributions of income and assets to populate the value table. With this in hand, at time point t , the goal is to run a deterministic DP in Line 4 using the precalculated table V and use that to calculate the optimal consumption at that time point t . This gives us an online algorithm that yields a new consumption at each time point t , based on the set of historical choices as well as the lookahead information available.

4 Theoretical Analysis

In this section, we establish a few theoretical results on the consumption behavior and the “power of lookahead”. We analyze the advantage that a worker privileged with lookahead, could attain compared to an underprivileged worker without lookahead. Let y_t, a_t, c_t be the income, assets and consumption at time t . Let T be the timeframe of the job and $u(c) = \sqrt{c}$ be the utility gained from consuming c .

Our first result shows that even in the very simple setting of $\beta = 1$ (no discount factor), and income y_t in the range $[0, Y]$, the difference between the total utility of algorithms with a k -step lookahead and without lookahead can be $\Omega(k\sqrt{Y})$. This is true not only of the algorithms we study, but *any* algorithm. This indicates that there are instances where a worker with lookahead privileges gains an edge over an underprivileged worker without lookahead, and the advantage grows linearly with the level of lookahead.

Theorem 1. *Consider two individuals, one with a lookahead of k steps and one with no lookahead. Let c_1, c_2, \dots, c_T be the consumption of the individual with lookahead k and z_1, z_2, \dots, z_T be the consumption of the individual with no lookahead. There exist instances where all incomes are in the range $[0, Y]$, such that*

$$\sum_{t=1}^T \sqrt{c_t} - \sum_{t=1}^T \sqrt{z_t} \geq \Omega(k\sqrt{Y}) \quad (4)$$

While our lower bound result is strong for large values of k , it is not very strong for small values. This is partly because we want to emphasize that a gap that grows with k holds even with $R_t = \beta = 1$.

We demonstrate that for any lookahead k , there exists an income distribution (dependent on k) such that the agent with lookahead achieves a $\Omega(k)$ advantage in utility over the agent without lookahead. This dependence on k does not weaken the claim, as our goal is merely to establish the existence of such instances.

Proof. As discussed above, we consider a very simple setting: $\beta = 1$, returns $R_t = 1$ for all t . Further we will assume that the individuals start with $a_1 = 0$ (no initial assets).

We consider the following input. Let Y be any parameter > 0 . Let $y_t = Y$ for $t \leq k/2$ and $y_t = x \cdot Y$ for $k/2 < t \leq k$ where x is a value uniformly sampled from $[0, 1]$. Note that both the individuals (the one with and without lookahead) know this input distribution. For simplicity, we also assume that the total time horizon T equals k (this assumption can be easily removed by setting $y_t = \frac{(1+x)}{2}Y$ for all $t > k$). As a final simplification, since the incomes can all be scaled, we can assume that $Y = 1$ for the remainder of the proof.

First, consider the individual with k lookahead. Note that they can see the value of x , and thus they can consume an amount $\frac{(1+x)}{2}$ at every time step. This is feasible because the first $k/2$ steps have income $y_t = 1$, and so the assets remain above the consumption at all time steps. This yields a total utility of $k\sqrt{\frac{1+x}{2}}$.

Now, consider an individual who does not have any lookahead. Intuitively, they cannot guess the value of x , and thus consuming $\frac{1+x}{2}$ is infeasible. But note that the individual may use a complex (possibly randomized) algorithm that consumes non-uniformly and achieve a high utility. We show that this is not possible.

The starting point of the proof is the classic minmax theorem of Yao (Motwani & Raghavan, 2013): for a given input distribution, an optimal algorithm for inputs drawn from this distribution, is a deterministic algorithm. In other words, in order to prove our desired lower bound, it suffices to restrict ourselves to deterministic algorithms and prove a bound on the difference in total utility, in expectation over the choice of x . For any deterministic algorithm, since in time steps $1, \dots, (k/2)$, the algorithm only sees income of 1, the values consumption $z_1, z_2, \dots, z_{(k/2)}$ are all fixed. Let $S = z_1 + z_2 + \dots + z_{(k/2)}$.

First, suppose it so happens that

$$\left| S - \frac{k}{2} \cdot \frac{1+x}{2} \right| > c \cdot k, \quad (5)$$

for some parameter c . In this case, we will argue that $\sum_i \sqrt{z_i}$ is significantly smaller than $k\sqrt{\frac{1+x}{2}}$. The starting point is the following inequality about the strict concavity of the square root function:

Lemma 1. *Let $a \in (1/2, 1)$ be a constant, and let $w \in (0, 1)$. Then we have:*

$$\sqrt{w} \leq \sqrt{a} + \frac{1}{2\sqrt{a}}(w - a) - \frac{1}{8}(w - a)^2. \quad (6)$$

The proof follows by a simple calculation.

Proof. We have:

$$\sqrt{w} - \sqrt{a} - \frac{1}{2\sqrt{a}}(w - a) = (w - a) \left(\frac{1}{\sqrt{w} + \sqrt{a}} - \frac{1}{2\sqrt{a}} \right) = \frac{(w - a)(a - w)}{2\sqrt{a}(\sqrt{w} + \sqrt{a})^2} = -\frac{(w - a)^2}{2\sqrt{a}(\sqrt{w} + \sqrt{a})^2} \quad (7)$$

The denominator is ≤ 8 , and thus by rearranging, the inequality follows. \square

Now, let us write $a = \frac{1+x}{2}$. By assumption, we have that $|z_1 + z_2 + \dots + z_{k/2} - \frac{k}{2}a| > ck$, and thus we have

$$\sum_{i \leq k} |z_i - a| > ck. \quad (8)$$

Next, we can use Lemma 1 to conclude that

$$\sum_{i \leq k} \sqrt{z_i} \leq k\sqrt{a} + \frac{1}{2\sqrt{a}} \sum_{i \leq k} (z_i - a) - \frac{1}{8} \sum_{i \leq k} (z_i - a)^2. \quad (9)$$

Now, since the consumption cannot be larger than the overall income (which is ka), the middle term on the RHS is ≤ 0 . Thus, we have

$$\sum_{i \leq k} \sqrt{z_i} \leq k\sqrt{a} - \frac{1}{8}(z_i - a)^2. \quad (10)$$

Next, by the Cauchy-Schwartz inequality and equation 8,

$$\sum_i (z_i - a)^2 \geq \frac{1}{k} \left(\sum_i |z_i - a| \right)^2 > c^2 k. \quad (11)$$

Together, the above inequalities imply that $\sum_i \sqrt{z_i} \leq k\sqrt{a} - \frac{c^2 k}{8}$. This shows that if the values z_i chosen by the deterministic algorithm satisfy equation 5, then the algorithm with no lookahead has total utility $\Omega(k)$ worse than an algorithm with lookahead.

The final step is to prove that if x is chosen at random from $(0, 1)$, the condition equation 5 holds with a constant probability for some $c > 0$. Since S is fixed, the condition is equivalent to $|\frac{2S}{k} - \frac{1+x}{2}| > 2c$. Equivalently, $|\frac{4S}{k} - 1 - x| > 4c$. For any fixed α , if $x \sim_{\text{uar}} (0, 1)$ the probability that $|\alpha - x| \leq 1/3$ is clearly $\leq 2/3$. Thus, the condition above must hold with $c = 1/12$, with probability at least $1/3$.

Therefore, we have that with probability $1/3$, the no-lookahead algorithm is $\Omega(k)$ worse than the algorithm with lookahead, and it can never be better. Thus the *expected* difference between the total utilities is also $\Omega(k)$. Yao's minmax theorem implies that this bound holds for any (possibly randomized) algorithm. \square

Theorem 1 shows that an algorithm with lookahead has a provable advantage over an algorithm that knows only the distribution of the incomes, even in the simplest setting where the decay factor $\beta = 1$ and the returns are all 1. Furthermore, the advantage grows *linearly* with the amount of lookahead. In particular, if an individual has infinite lookahead, they can have an advantage of $\Omega(T)$.

Next, we show that when $\beta = 1$ and $R_t = 1$ for all t , the lower bound from Theorem 1 cannot be improved.

Lemma 2. *Suppose $\beta = 1$ and the return $R_t = 1$ for all t . Let y_1, y_2, \dots, y_T be a sequence of income values with $y_t \geq 0$ for all t . Then if the consumption sequence of a k -lookahead algorithm is $\{c_1, c_2, \dots, c_T\}$, there exists a no-lookahead algorithm whose total utility is $\sqrt{c_1} + \sqrt{c_2} + \dots + \sqrt{c_{T-k}}$.*

In other words, the difference in the total utility (between the k -lookahead and the no-lookahead algorithms) is simply $\sqrt{c_{T-k+1}} + \dots + \sqrt{c_T}$. In settings where all the c_t are of magnitude $O(\text{income})$, this corresponds to $O(k)$ times the square root of the income, which is the lower bound in Theorem 1.

Proof. The proof is simple: a no-lookahead algorithm can mimic a k -lookahead algorithm, but with a delay of k steps. We will call this the k -delay algorithm. Let c_1, c_2, \dots, c_T be the consumption sequence of the given k -lookahead algorithm. The k -delay algorithm is defined as follows,

- For $t = 1$ to T :
1. If $t \leq k$, consume 0
 2. Else consume c_{t-k}

The total utility bound is easy to see. One only needs to check that the algorithm is feasible (i.e., it satisfies the condition that the total consumption until time t is bounded by the total income plus the assets until that time). This is easy to check because the algorithm consumes 0 for the first k steps, while the income $y_t \geq 0$. Since the decay factor $\beta = 1$, delay does not change the utility the algorithm receives. \square

Remark. We see that the proof relies on having $\beta = 1$. If we take into account factors such as inflation (e.g., $\beta = 0.95$), then the “power of lookahead” can likely be made much more significant. The agents with lookahead could potentially use the lookahead information to anticipate future gains and consume more than the expected consumption and therefore gain a significant advantage over the agents without lookahead.

5 Experiments

In this section, using our online learning algorithm, we explore three different sets of experiments regarding the effects of lookahead, the parameters involved in lookahead under uncertain job timetables, and mitigation schemes.

5.1 Experiment Setup

We will first introduce the different elements in our experiment setup.²

Workers. Within our framework, agents represent employed individuals who earn weekly income, own assets, and decide whether to consume or save. We create an income distribution of 10,000 individuals using 2019 income data from the US Census Bureau’s Annual ASEC survey of the Consumer Price Index (by the IPUMS Consumer Price Survey) (Flood et al., 2020; PK, 2019-2020). In the first stage, we eliminate outliers from the income distribution to address challenges related to individuals with exceptionally high or low earnings, which can be challenging to compare due to significant scale differences. We identify outliers by employing the commonly-used inter-quartile range (IQR) proximity rule (Dekking et al., 2005). Subsequently, we categorize individuals into four distinct income groups: low incomes ranging from \$1.22 to \$1125.49, low-middle incomes from \$1125.49 to \$2249.75, high-middle incomes from \$2249.75 to \$3374.02, and high incomes from \$3374.02 to \$4498.29. This classification is achieved by partitioning the overall income range into these four segments. To establish asset values, individuals are assigned the median population asset value of \$123,840, derived from median percentile net worth data and median net worth by income percentile data from the Federal Reserve (of Governors of the Federal Reserve System, 2019).

Minimum Subsistence. Furthermore, our simulation considers minimum subsistence, i.e., the constraint that the individuals must allocate funds for their minimum basic needs, such as food and shelter (Zhang & Liu, 2020; Zimmerman & Carter, 2003; Álvarez Peláez & Díaz, 2005; Shin & Lim, 2011; Shim & Shin, 2014; Antony & Klarl, 2019; Dwork & Ilvento, 2018; Miranda-Pinto et al., 2023; 2020). This consideration acknowledges the fundamental requirement for individuals to fulfill their basic necessities as part of the decision-making process regarding consumption and utility maximization. The minimum subsistence values are derived from mean annual expenditures in 2019 from the Consumer Expenditure Surveys of the US Bureau Of Labor Statistics (of Labor Statistics, 2019), stratified based on income levels. In essence, individuals with similar income levels are obligated to cover equivalent amounts for their basic needs.

Shocks. Shocks, defined as alterations to an agent’s financial state due to schedule instability, play a pivotal role in influencing decisions related to consumption and savings. Shocks can either positively or negatively impact income, such as sudden work-hour reductions or increases. The magnitude of income shocks ranges from -0.4 to 0.4 , with shocks occurring as $(1 + r) \times \text{income}$, where r represents the shock value and the income represents the current value of income uniformly sampled from the income distribution. The shocks are generated from a Bernoulli process, with the shock size parameter r uniformly sampled from $[-0.4, 0.4]$. We set the probability of shocks to 0.4 .

Isolating the Impact of Lookahead and Other Parameters. At each time point (here, a workweek), each worker within an income group shares identical observable features, like income, shocks stemming from unpredictable schedules, and returns on assets. The sole point of divergence among individuals within the same group is the extent of lookahead they possess. This deliberate design choice aims to isolate the temporal impact of lookahead on utility, eliminating other financial factors’ interference. The overall job timeline spans 26 weeks, i.e., a 6-month job duration. The returns are sampled from uniformly from $[0.9, 1.1]$ and with probability 0.1 there is an additive shock uniformly sampled from $[-0.05, 0.05]$ on the return rates. The discounting factor β is set to the commonly-used value of 0.95 .

²The anonymous code can be accessed at https://osf.io/rkx3v/?view_only=a35d6ecd9ff7488e94766db333a16f22

Further Clarification on Simulation Parameters. Supplementary explanations regarding the simulation parameters (along with their grounding in real-world data and statistics) are provided in Appendix §A.2.

5.2 Future Lookahead: An Empirical Inquiry

This section aims to examine the utility acquired by individuals under varying levels of lookahead. Specifically, the experiment seeks to compare those with minimal (or no) information about the future, such as protected groups like hourly, part-time workers, and specific racial groups that typically receive limited or no advance notice, against other workers with more foresight. The experiment setup is as explained in §5.1 and the outcomes of this experiment are depicted in Figure 1.




Figure 1: The final utility gained for different levels of lookahead is illustrated for four distinct income classes, each comprising 27 individuals. Workers are of similar features, with variations solely based on the temporal aspect, i.e., the amount of lookahead in their work schedules. The x -axis depicts the lookahead value and the y -axis represents the total utility at the end of T steps. The two orange triangles are in place to highlight the difference in the level of lookahead needed for approaching the near-maximum utility (near-maximum utility is a utility close to that of a worker with full information at L26) as the income levels change.

Analysis. The insights from Figure 1 can be summarized in four key points. 1) Firstly, workers with lower lookahead and minimal information about future instability experience significantly lower financial utility compared to those with higher lookahead. Unsurprisingly, individuals struggle to manage their finances effectively when confronted with unforeseen schedules. 2) Secondly, lookahead has a positive impact overall. Workers with more lookahead can efficiently manage their finances, resulting in higher utility. 3) Thirdly, workers do not require full information about their work schedules in advance. Even beyond the midpoint lookahead (lookahead 13), workers can achieve a utility comparable to those with complete information about their schedules. 4) Lastly, higher income leads to greater utility, as individuals can consume without concerns. This is evident from the consistent shift of the plots along the y -axis.

However, irrespective of income level, having more than the midpoint level of lookahead proves advantageous for individuals. Providing people with advance notice of their schedules can be reasonably implemented for the next 2-3 months of work without requiring employers to furnish complete information at the start of their tenure. Notably, for higher-income workers, less future information is needed to approach the utility values of someone with complete information, as observed by the leftward movement of the orange triangles across all lookahead levels.

5.3 Dynamics of Asset Appreciation and Depreciation

Our simulated space provides a comprehensive platform for delving into the intricacies of parameters within work scheduling. This exploration includes understanding the behavior of individuals with diminishing assets compared to those experiencing favorable returns on their savings.

Thus, this section’s objective is to investigate the impact of asset appreciation and depreciation, i.e., positive and negative return rates on workers’ assets on decision-making across various levels of lookahead. This also serves as an examination of scenarios wherein the workers are already at a (dis)advantage in terms of assets.

Assets depreciate when their value declines over time, influenced by factors like risky investments, fluctuating market conditions, tax obligations, possessions becoming obsolete, and wear and tear, as seen in vehicles, buildings, and cash saved without earning interest. Conversely, asset appreciation occurs when individuals receive returns on their savings or make profitable investments in stocks, among other factors.

All experimental settings are similar to that of §5.1 other than the return rates on assets. Here, we compare negative return rates in the range $[0.75, 0.95]$ to positive return rates in $[1.05, 1.25]$.

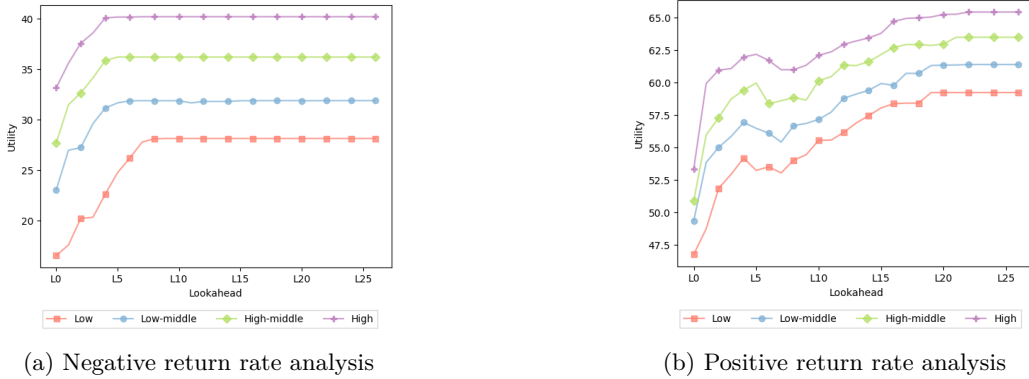


Figure 2: Individuals with similar features but varying levels of lookahead under negative return rates ranging from 0.75 to 0.95 on their assets, as well as positive return rates between 1.05 and 1.25 on their assets.

Analysis. The findings are illustrated in Figure 2. Several insights emerge from this analysis: 1) Firstly, workers generally experience higher utility when they encounter favorable return rates. This is evident by comparing the utility range in the negative rates plot, which spans from 15 to 40, to the positive plot, which ranges from 45 to 65. 2) Secondly, with depreciating assets, individuals benefit significantly from small amounts of lookahead, reaching near-maximum utility. People across income classes achieve a near-maximum utility before Lookahead 10. There is a consistent decline in the lookahead value required to attain near-maximum utility, reaching around lookahead 5 for the highest income class and around 10 for the lowest income class. This observation aligns with the trend observed in the previous section, where more income classes require less future information to reach peak utility values. 3) Thirdly, in the case of positive returns, individuals have more flexibility in consumption, as they anticipate overall favorable returns ahead. Future information is not as crucial as in the negative returns scenario, where they are not at a disadvantageous situation with depreciating assets. This explains the late convergence of all income classes to the peak utility value (under positive return rates).

5.4 Interventions

In this section, the goal is to examine the mitigating effects of two intervention scenarios. In terms of intervention policies, simulation provides a valuable sandbox environment for testing that might be challenging or even impossible to explore in the real world. In a simulated work scheduling setting, researchers, policymakers, and practitioners can experiment with various interventions, assess their effects, and fine-tune strategies without real-world consequences.

All experimental settings are similar to §5.1. We assign interventions (to a sample of 50 individuals per intervention), as follows: **(1) Compensation:** Workers will be compensated for sudden schedule changes and *on-call* shifts. This policy is modeled after the measures approved by the San Francisco Board of Supervisors, which introduced new protections for retail workers in the city, necessitating employers to provide compensation for unpredictable schedules based on factors such as employment type, hourly rate,

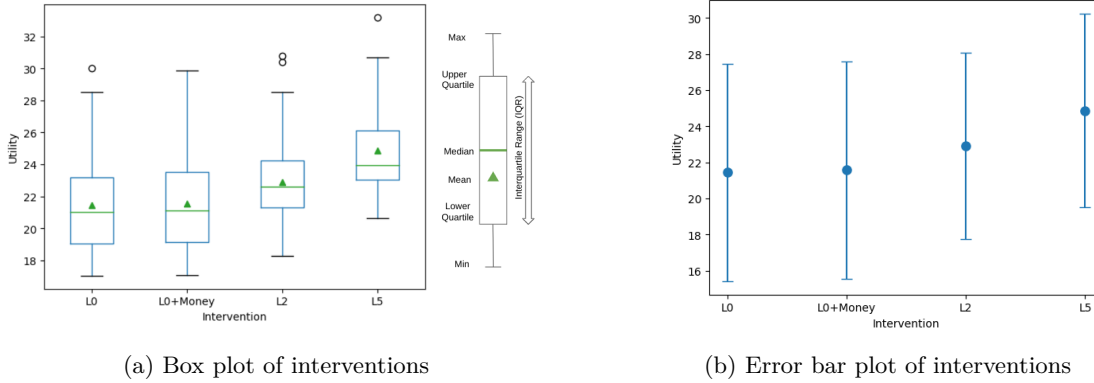


Figure 3: The box plot and the error bar plot of the statistical distribution of additional gained utilities for various interventions. The interventions considered are L0 + Money, which involves compensation fees for individuals with no future information; L2, assigning at least two weeks of lookahead to all; and L5, assigning at least five weeks of lookahead to everyone. In the box plot, the green triangle represents the mean, and the green line represents the median. In the error bar plot, the blue circle represents the mean and the error bars are 2 standard deviations from the mean (95% confidence interval).

hours of work (Golden, 2015). Inspired by this policy, we disburse twice the amount of earnings back when there is a negative shock. **(2) Mandatory minimum advance notice:** Every worker should be entitled to a mandated minimum lookahead, i.e., they should be aware of their schedule for the upcoming two weeks. This proposition is inspired by the Schedules that Work Act of 2014 (H.R. 5159) presented in Congress, which stipulates that if there are alterations to the schedule and minimum hours, the employer must inform the employee at least two weeks prior to the start point of the new schedule (Golden, 2015). Inspired by this policy, we assign workers two weeks and roughly one month (5 weeks) of minimum lookahead, separately.

Analysis. Results for intervention scenarios are depicted in Figure 3 as a box plot to show the distributions (and a separate corresponding error bar plot to capture uncertainty). Three insights can be derived: 1) Firstly, interventions have a positive impact overall, evident from the increased utility across all statistical metrics (mean, median, confidence interval, and all quartiles) post-intervention. 2) Secondly, while compensation in the form of additional income for unanticipated schedule changes is beneficial, it does not substitute for providing workers with advance notice of their shift schedules. Having knowledge of future plans with the same income but a predictable schedule appears to be more effective for workers’ financial well-being than an unforeseen, volatile schedule compensated with fees. A comparison between *L0 + Money* and *L2* indicates that even incorporating two weeks of lookahead is more advantageous than providing compensation for instability without any lookahead. 3) Thirdly, offering just one month of advance notice results in a notable increase in utility, as seen in the comparison between *L5* and *L0*. Even small amounts of notice can significantly enhance utility.

Note on the Robustness of Experimental Parameters/Results. We have included the results for 20 different randomly generated seeds (with a 95% confidence interval) in Appendix §A.1 in Figure 4. It is important to mention that we conducted these experiments with various random seeds, considering different runs (e.g., with a median asset of approximately \$70,000 (Aurelia Glass, 2023), representing the median working-class wealth, as opposed to the population median), and with other realistic discount and return parameters, as well as varied plausible shock sizes. The overall trends in our results in the previous sections remain consistent as long as the chosen parameters fall within more realistic ranges. If one opts for more extreme and unrealistic parameters, the distinctions observed in §5.2, 5.3, and 5.4 will become even more pronounced.

Note on Additional Real-world Scenarios. To explore more diverse settings, adopt our framework/simulation in more real-world problems, and extend our framework to more realistic contexts, we incorporate three additional case studies: the role of schedule instability during unemployment in §B.2, the

influence of educational attainment on consumption behavior and patience for future information in §B.3, and a more granular analysis of return rates by comparing cash and stocks in §B.4. Across these diverse scenarios, we observe behavioral patterns consistent with those identified in the core sections. Further details can be found in the corresponding appendix sections.

6 Discussions and Limitations

The primary contribution of our paper lies in the development of an online framework that delves into the financial insecurity stemming from work schedule instability. This framework addresses two crucial questions pertinent to work scheduling: 1) How can we develop an online learning approach for consumption under uncertainty to formally and empirically assess the negative effects of work schedule instability on income/employment? 2) How can this in turn help with informing the development of effective mitigation policies and regulations in the workplace? We provide analytical insights into how lookahead significantly enhances individuals’ utility, with a focus on the effectiveness of increased lookahead. Our empirical investigation employs simulations and explores two distinct intervention strategies aimed at mitigating the adverse effects of schedule instability. While our model is being deployed in a work setting, it applies to any scenario that involves temporal uncertainty with the possibility of improvement with lookahead. Our paper, however, has a number of limitations:

Simulation. A simulation is only as good as the models used to build it and relies on idealized and formal models of human behavior that are simplifications. A simplified model however is useful for an inquiry into *limits*: in our work, we show that even under ideal models of human utility-maximizing behavior, the lack of predictability and lookahead has concrete consequences for financial stability.

Homogeneity. The simulation assumes societal uniformity within income groups, considering temporal differences as the main variation. However, this overlooks other social inequities and heterogeneities that could impact decision-making. Income instability can be influenced by factors such as gender (Blau & Kahn, 2017), ability (Whittaker et al., 2019), race (Zwerling & Silver, 1992), and health status (Hicken et al., 2014). Additionally, access to public policies addressing these disparities may also differ across demographic groups.

On Algorithmic Scheduling. In the digital age, algorithmic processes and digital technologies play a crucial role across various workforce sectors, from business process management to automated work scheduling. These tools enhance efficiency and productivity, offering advantages over manual scheduling methods by saving time and effort while providing greater control to managers and employees. However, despite these benefits, algorithmic scheduling can also introduce new challenges (Lancaster, 2015).

Scheduling analytics have the potential to enable operators to exploit their workforce unfairly. Reports from The New York Times point to scheduling software as a factor behind more unpredictable scheduling practices, particularly for low-wage workers in the service industry (Lancaster, 2015). From the standpoint of employees, this unpredictability in scheduling is perceived as a negative attribute. It is an instability enforced by employers, often utilizing workforce management technology and algorithms. Therefore, it would be beneficial to enhance our modeling approach by incorporating more provisions to account for the intricate details and mechanisms of algorithmic scheduling parameters and operations.

Future Directions. Enhancing this framework could involve integrating additional financial factors linked to temporal uncertainty, such as retirement savings constraints, differentiation between risky and riskless assets, and broader socioeconomic shifts like workplace restructuring or job relocations. Additionally, the framework offers opportunities to explore fairness and inequity in automated shift scheduling, particularly for sub-populations with sensitive attributes. It also enables research into targeted interventions within specific work environments (e.g., food and retail sectors) to mitigate the long-term effects of biased decision-making.

References

- Mitsuteru Abe, Fabio Tanaka, Jair Pereira Junior, Anna Bogdanova, Tetsuya Sakurai, and Claus Aranha. Using agent-based simulator to assess interventions against covid-19 in a small community generated from map data. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '22, pp. 1–8, Richland, SC, 2022. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450392136.
- Jürgen Antony and Torben Klarl. Non-renewable resources, subsistence consumption, and hartwick’s investment rule. *Resource and Energy Economics*, 55:124–142, 2019. ISSN 0928-7655. doi: <https://doi.org/10.1016/j.reseneeco.2018.11.002>. URL <https://www.sciencedirect.com/science/article/pii/S0928765518300290>.
- Christian E. Weller Aurelia Glass, David Madland. Unions build wealth for the american working class, 2023. URL <https://www.americanprogress.org/article/unions-build-wealth-for-the-american-working-class/#:~:text=A%202022%20survey%20found%20that,in%20that%20period%2C%20and%2011.4>.
- Neil Bania and Laura Leete. Income volatility and food insufficiency in us low-income households, 1991-2003. Technical report, Citeseer, 2006.
- Deutsche Bank. The age of disorder, 2020. <http://www.epge.fr/wp-content/uploads/2020/09/The-age-of-disorder.pdf>.
- Erhan Bayraktar and Virginia R. Young. Maximizing utility of consumption subject to a constraint on the probability of lifetime ruin, 2012. URL <https://arxiv.org/abs/1206.6268>.
- Jess Benhabib and Alberto Bisin. Skewed wealth distributions: Theory and empirics. *Journal of Economic Literature*, 56(4):1261–91, 2018.
- Francine D Blau and Lawrence M Kahn. The gender wage gap: Extent, trends, and explanations. *Journal of economic literature*, 55(3):789–865, 2017.
- Monika Büttler. Neoclassical life-cycle consumption: A textbook example. *Economic Theory*, 17(1):209–221, 2001.
- Ben Carlson. 200 years of asset class returns, 2021. <https://awealthofcommonsense.com/2021/05/200-years-of-asset-class-returns/>.
- Christopher F Chabris, David I Laibson, and Jonathon P Schuldt. Intertemporal choice. In *Behavioural and experimental economics*, pp. 168–177. Springer, London, 2010.
- CLASP. California’s low-wage retail workers want off the roller coaster of unpredictable schedules, 2022. URL <https://www.clasp.org/blog/californias-low-wage-retail-workers-want-roller-coaster-unpredictable-schedules/>.
- Russell Cooper and Jonathan Willis. Discounting: Investment sensitivity and aggregate implications. *Unpublished manuscript*, 2014.
- M. Cristelli, L. Pietronero, and A. Zaccaria. Critical overview of agent-based models for economics, 2011.
- Gal Dalal, Assaf Hallak, Steven Dalton, Shie Mannor, Gal Chechik, et al. Improve agents without retraining: Parallel tree search with off-policy correction. *Advances in Neural Information Processing Systems*, 34: 5518–5530, 2021.
- Angus Deaton. Saving and liquidity constraints. Technical report, National Bureau of Economic Research, 1989.
- Angus Deaton et al. *Understanding consumption*. Oxford University Press, Oxford, 1992.

- Frederik Michel Dekking, Cornelis Kraaikamp, Hendrik Paul Lopuhaä, and Ludolf Erwin Meester. *A Modern Introduction to Probability and Statistics: Understanding why and how*, volume 488. Springer, London, 2005.
- Wouter J Den Haan. Comparison of solutions to the incomplete markets model with aggregate uncertainty. *Journal of Economic Dynamics and Control*, 34(1):4–27, 2010.
- Robert Dochow. *Online algorithms for the portfolio selection problem*. Springer, Wiesbaden, Germany, 2016.
- Cynthia Dwork and Christina Ilvento. Fairness under composition. In *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*, Wadern/Saarbruecken, Germany, 2018. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik GmbH.
- Alexander D’Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, D. Sculley, and Yoni Halpern. Fairness is not static: Deeper understanding of long term fairness via simulation studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* ’20*, pp. 525–534, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450369367. doi: 10.1145/3351095.3372878. URL <https://doi.org/10.1145/3351095.3372878>.
- Yonathan Efroni, Gal Dalal, Bruno Scherrer, and Shie Mannor. Beyond the one-step greedy approach in reinforcement learning. In *International Conference on Machine Learning*, pp. 1387–1396, Stockholm SMässa, Stockholm SWEDEN, 2018. PMLR, PMLR.
- Yonathan Efroni, Mohammad Ghavamzadeh, and Shie Mannor. Online planning with lookahead policies. *Advances in Neural Information Processing Systems*, 33:14024–14033, 2020.
- Hadi Elzayn, Shahin Jabbari, Christopher Jung, Michael Kearns, Seth Neel, Aaron Roth, and Zachary Schutzman. Fair algorithms for learning in allocation problems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT* ’19*, pp. 170–179, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450361255. doi: 10.1145/3287560.3287571. URL <https://doi.org/10.1145/3287560.3287571>.
- Danielle Ensign, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. Runaway feedback loops in predictive policing. In Sorelle A. Friedler and Christo Wilson (eds.), *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, pp. 160–171, New York, NY, USA, 23–24 Feb 2018. PMLR. URL <https://proceedings.mlr.press/v81/ensign18a.html>.
- Keith Marzilli Ericson and David Laibson. Intertemporal choice. In *Handbook of behavioral economics: Applications and foundations 1*, volume 2, pp. 1–67. Elsevier, 2019.
- Diana Farrell and Fiona Greig. Paychecks, payday, and the online platform economy. In *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association*, volume 109, pp. 1–40, Washington, DC, USA, 2016. JSTOR, National Tax Association.
- Sarah Flood, Miriam King, Renae Rodgers, Steven Ruggles, and J Robert Warren. Integrated public use microdata series, current population survey: version 7.0 [dataset]. *Minneapolis, MN: IPUMS*, 10:D030, 2020.
- Milton Friedman. *Theory of the consumption function*. Princeton university press, Princeton, New Jersey, 2018.
- Peter Ganong and Pascal Noel. Consumer spending during unemployment: Positive and normative implications. *American economic review*, 109(7):2383–2424, 2019.
- Lisa A Gennetian and Eldar Shafir. The persistence of poverty in the context of financial instability: A behavioral perspective. *Journal of Policy Analysis and Management*, 34(4):904–936, 2015.
- Lonnie Golden. Irregular work scheduling and its consequences, 2015.

- Peter Grandits. An optimal consumption problem in finite time with a constraint on the ruin probability. *Finance and Stochastics*, 19(4):791–847, 2015.
- Kathleen Griesbach, Adam Reich, Luke Elliott-Negri, and Ruth Milkman. Algorithmic control in platform food delivery work. *Socius*, 5:2378023119870041, 2019. doi: 10.1177/2378023119870041. URL <https://doi.org/10.1177/2378023119870041>.
- Anthony Hannagan and Jonathan Morduch. Income gains and month-to-month income volatility: Household evidence from the us financial diaries. *SSRN Electronic Journal*, Online:Paper No. 2659883, 21, 2015. ISSN 1556-5068. doi: 10.2139/ssrn.2659883.
- Colin B Hawkinson, Sarah B Andrea, Anjum Hajat, Anita Minh, Shanise Owens, Kieran Blaikie, Jessie Seiler, Andrea R Molino, and Vanessa M Oddo. A cross-sectional analysis of work schedule notice and depressive symptoms in the united states. *SSM Popul. Health*, 22(101413):101413, June 2023.
- Julia R. Henly and Susan J. Lambert. Unpredictable work timing in retail jobs: Implications for employee work–life conflict. *ILR Review*, 67(3):986–1016, 2014. doi: 10.1177/0019793914537458. URL <https://doi.org/10.1177/0019793914537458>.
- Margaret T Hicken, Hedwig Lee, Jeffrey Morenoff, James S House, and David R Williams. Racial/ethnic disparities in hypertension prevalence: reconsidering the role of chronic stress. *American journal of public health*, 104(1):117–123, 2014.
- Lily Hu, Nicole Immorlica, and Jennifer Wortman Vaughan. The disparate effects of strategic manipulation. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* ’19, pp. 259–268, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450361255. doi: 10.1145/3287560.3287597. URL <https://doi.org/10.1145/3287560.3287597>.
- Mia Hubert and Stephan Van der Veeken. Outlier detection for skewed data. *Journal of Chemometrics: A Journal of the Chemometrics Society*, 22(3-4):235–246, 2008.
- Self Inc. What happens if you get a job while on unemployment?, 2021. URL <https://www.self.inc/blog/when-happens-if-you-get-job-on-unemployment>.
- Business Insider. Uber, lyft, instacart gig workers saw earnings fall — here’s how much they made in 2024, 2024. URL <https://www.businessinsider.com/uber-lyft-instacart-gig-workers-saw-earnings-fall-2024-2025-2>.
- Accounting Insights. Analyzing right-skewed distributions in financial data. *Accounting Insights*, 2025. URL <https://accountinginsights.org/analyzing-right-skewed-distributions-in-financial-data/#::~text=By%20focusing%20on%20the%20range%20between%20the,understanding%20the%20variability%20in%20right%20skewed%20data>.
- Brookings Institution. How does unemployment insurance work—and how is it changing during the coronavirus pandemic?, 2020. URL <https://www.brookings.edu/articles/how-does-unemployment-insurance-work-and-how-is-it-changing-during-the-coronavirus-pandemic>.
- JPMorgan Chase Institute. Weathering volatility: Big data on the financial ups and downs of u.s. individuals, 2016. URL <https://www.jpmorganchase.com/content/dam/jpmc/jpmorgan-chase-and-co/institute/pdf/institute-volatility-cash-buffer-report.pdf>.
- Arne L Kalleberg. Precarious work, insecure workers: Employment relations in transition. *American sociological review*, 74(1):1–22, 2009.
- Arne L Kalleberg. *Good jobs, bad jobs: The rise of polarized and precarious employment systems in the United States, 1970s-2000s*. Russell Sage Foundation, New York, USA, 2011.
- Sampath Kannan, Aaron Roth, and Juba Ziani. Downstream effects of affirmative action. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* ’19, pp. 240–248, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450361255. doi: 10.1145/3287560.3287578. URL <https://doi.org/10.1145/3287560.3287578>.

- Jodi Kantor. Working anything but 9 to 5. *The New York Times*, 13, 2014.
- Franziska Klügl and Hildegunn Kyvik Nordås. Modelling agent decision making in agent-based simulation – analysis using an economic technology uptake model. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS ’23, pp. 1903–1911, Richland, SC, 2023. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450394321.
- Moritz Kuhn. Recursive equilibria in an aiyagari-style economy with permanent income shocks. *International Economic Review*, 54(3):807–835, 2013.
- David Laibson, Sean Chanwook Lee, Peter Maxted, Andrea Repetto, and Jeremy Tobacman. Estimating discount functions with consumption choices over the lifecycle. *The Review of Financial Studies*, pp. hhae035, 2024.
- Susan J Lambert. Passing the buck: Labor flexibility practices that transfer risk onto hourly workers. *Human relations*, 61(9):1203–1227, 2008.
- Susan J. Lambert, Julia R. Henly, and Jaeseung Kim. Precarious work schedules as a source of economic insecurity and institutional distrust. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 5(4):218–257, 2019. ISSN 2377-8253. doi: 10.7758/RSF.2019.5.4.08. URL <https://www.rsfsjournal.org/content/5/4/218>.
- Alaina Lancaster. The pros and cons of scheduling software, 2015. URL <https://www.restaurantbusinessonline.com/technology/pros-cons-scheduling-software>.
- Min Kyung Lee, Ishan Nigam, Angie Zhang, Joel Afriyie, Zhizhen Qin, and Sicun Gao. Participatory algorithmic management: Elicitation methods for worker well-being models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’21, pp. 715–726, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384735. doi: 10.1145/3461702.3462628. URL <https://doi.org/10.1145/3461702.3462628>.
- Laura Leete and Neil Bania. The effect of income shocks on food insufficiency. *Review of Economics of the Household*, 8:505–526, 2010.
- Bin Li and Steven CH Hoi. Online portfolio selection: A survey. *ACM Computing Surveys (CSUR)*, 46(3): 1–36, 2014.
- Bin Li and Steven Chu Hong Hoi. *Online portfolio selection: principles and algorithms*. Crc Press, Boca Raton, 2018.
- Lydia T. Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. Delayed impact of fair machine learning. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3150–3158, Stockholm"assan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL <https://proceedings.mlr.press/v80/liu18c.html>.
- Benjamin B Lockwood. Optimal income taxation with present bias. *American Economic Journal: Economic Policy*, 12(4):298–327, 2020.
- K Loggins. Here’s what happens when an algorithm determines your work schedule, 2020.
- Qingyin Ma, John Stachurski, and Alexis Akira Toda. The income fluctuation problem and the evolution of wealth. *Journal of Economic Theory*, 187:105003, 2020.
- Bill McCarthy, Angela Carter, Mikael Jansson, Cecilia Benoit, and Ryan Finnigan. Poverty, material hardship, and mental health among workers in three front-line service occupations. *Journal of Poverty*, 22(4): 334–354, 2018.
- Elaine D McCrate. Unstable and on-call work schedules in the united states and canada, 2018.

- Nadav Merlis. Reinforcement learning with lookahead information. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 64523–64581. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/7664a7e946a84ac5e97649a967717cf2-Paper-Conference.pdf.
- Nadav Merlis, Dorian Baudry, and Vianney Perchet. The value of reward lookahead in reinforcement learning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=URyeU8mwz1>.
- Jorge Miranda-Pinto, Daniel Murphy, Kieran James Walsh, and Eric R Young. A model of expenditure shocks, 2020.
- Jorge Miranda-Pinto, Daniel Murphy, Kieran James Walsh, and Eric R Young. Saving constraints, inequality, and the credit market response to fiscal stimulus. *European Economic Review*, 151:104355, 2023.
- Thomas M. Moerland, Anna Deichler, Simone Baldi, Joost Broekens, and Catholijn M. Jonker. Think too fast nor too slow: The computational trade-off between planning and reinforcement learning, 2020.
- Jonathan Morduch and Julie Siwicki. In and out of poverty: Episodic poverty and income volatility in the us financial diaries. *Social Service Review*, 91(3):390–421, 2017.
- Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, USA, 2013. ISBN 0511814070.
- Pegah Nokhiz, Aravinda Kanchana Ruwanpathirana, Neal Patwari, and Suresh Venkatasubramanian. Precarity: Modeling the long term effects of compounded decisions on individual instability. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’21, pp. 199–208, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384735. doi: 10.1145/3461702.3462529. URL <https://doi.org/10.1145/3461702.3462529>.
- Board of Governors of the Federal Reserve System. Changes in u.s. family finances from 2016 to 2019: Evidence from the survey of consumer finances, 2019. URL https://www.federalreserve.gov/econres/scf/dataviz/scf/table/#series:Net_Worth;demographic:inccat;population:all;units:median.
- U.S. Bureau of Labor Statistics. Consumer expenditure surveys, 2019. URL <https://www.bls.gov/cex/tables.htm>.
- Center on Budget and Policy Priorities. How many weeks of unemployment compensation are available?, 2021. URL <https://www.cbpp.org/research/economy/how-many-weeks-of-unemployment-compensation-are-available>.
- Jeffrey Parker. Theories of consumption and saving. *Economics*, 314:1–25, 2010. URL <https://www.reed.edu/economics/parker/s11/314/book/Ch16.pdf>.
- Arpita Patnaik, Joanna Venator, Matthew Wiswall, and Basit Zafar. The role of heterogeneous risk preferences, discount rates, and earnings expectations in college major choice. *Journal of Econometrics*, 231(1): 98–122, 2022.
- PK. Average, median, top 1%, and all united states household income percentiles in 2020, 2019-2020. URL <https://dqydj.com/average-median-top-net-worth-percentiles/>.
- Guillaume Rabault. When do borrowing constraints bind? some new results on the income fluctuation problem. *Journal of Economic Dynamics and Control*, 26(2):217–245, 2002.
- Federal Reserve. Report on the economic well-being of us households in 2015, 2016.
- Paul A Samuelson. A note on measurement of utility. *The review of economic studies*, 4(2):155–161, 1937.
- Thomas Sargent and John Stachurski. Quantitative economics. Technical report, Citeseer, 2014.

- Daniel Schneider and Kristen Harknett. Income volatility in the service sector: Contours, causes and consequences, 2017.
- Daniel Schneider and Kristen Harknett. Consequences of routine work-schedule instability for worker health and well-being, 2019a.
- Daniel Schneider and Kristen Harknett. It’s about time: How work schedule instability matters for workers, families, and racial inequality, 2019b.
- Anuj K Shah, Eldar Shafir, and Sendhil Mullainathan. Scarcity frames value. *Psychological science*, 26(4): 402–412, 2015.
- Gyoocheol Shim and Yong Hyun Shin. Portfolio selection with subsistence consumption constraints and cara utility. *Mathematical Problems in Engineering*, 2014:153793, Apr 2014. ISSN 1024-123X. doi: 10.1155/2014/153793. URL <https://doi.org/10.1155/2014/153793>.
- Yong Hyun Shin and Byung Hwa Lim. Comparison of optimal portfolios with and without subsistence consumption constraints. *Nonlinear Analysis*, 74(1):50–58, 2011. doi: 10.1016/j.na.2010.08.014.
- Harshit Sikchi, Wenxuan Zhou, and David Held. Learning off-policy with online planning. In *Proceedings of the 5th Conference on Robot Learning*, pp. 1622–1633, London, UK, 2022. PMLR, PMLR.
- Alarith Uhde, Nadine Schlicker, Dieter P. Wallach, and Marc Hassenzahl. Fairness and decision-making in collaborative shift scheduling systems. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI ’20, pp. 1–13, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450367080. doi: 10.1145/3313831.3376656. URL <https://doi.org/10.1145/3313831.3376656>.
- Christian E Weller. Working-class families are getting hit from all sides, 2018. URL <https://www.americanprogress.org/article/working-class-families-getting-hit-sides/>.
- Dewey Lonzo Whaley III. The interquartile range: Theory and estimation. Master’s thesis, East Tennessee State University, 2005.
- Meredith Whittaker, Meryl Alper, Cynthia L Bennett, Sara Hendren, Liz Kaziunas, Mara Mills, Meredith Ringel Morris, Joy Rankin, Emily Rogers, Marcel Salas, et al. Disability, bias, and ai. *AI Now Institute*, November, 2019.
- Alex J Wood. Powerful times: Flexible discipline and schedule gifts at work. *Work, Employment and Society*, 32(6):1061–1077, 2018.
- Alex J. Wood. Algorithmic Management: Consequences for Work Organisation and Working Conditions. JRC Working Papers on Labour, Education and Technology 2021-07, Joint Research Centre, May 2021. URL <https://ideas.repec.org/p/ipt/laedte/202107.html>.
- Wenzhi Xi, Zhanfeng Li, Xinyuan Song, and Hanwen Ning. Online portfolio selection with predictive instantaneous risk assessment. *Pattern Recognition*, 144:109872, 2023.
- Nutchanon Yongsatianchot, Noah Chicoine, Jacqueline Griffin, Ozlem Ergun, and Stacy Marsella. Agent-based modeling of human decision-makers under uncertain information during supply chain shortages. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS ’23, pp. 1886–1894, Richland, SC, 2023. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450394321.
- Angie Zhang, Alexander Boltz, Chun Wei Wang, and Min Kyung Lee. Algorithmic management reimaged for workers and by workers: Centering worker well-being in gig work. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI ’22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391573. doi: 10.1145/3491102.3501866. URL <https://doi.org/10.1145/3491102.3501866>.
- Xueru Zhang and Mingyan Liu. Fairness in learning-based sequential decision algorithms: A survey, 2020.

Stephan Zheng, Alexander Trott, Sunil Srinivasa, Nikhil Naik, Melvin Gruesbeck, David C Parkes, and Richard Socher. The ai economist: Improving equality and productivity with ai-driven tax policies, 2020.

Frederick J. Zimmerman and Michael R. Carter. Asset smoothing, consumption smoothing and the reproduction of inequality under risk and subsistence constraints. *Journal of Development Economics*, 71(2):233–260, 2003. ISSN 0304-3878. doi: [https://doi.org/10.1016/S0304-3878\(03\)00028-2](https://doi.org/10.1016/S0304-3878(03)00028-2). URL <https://www.sciencedirect.com/science/article/pii/S0304387803000282>.

Craig Zwerling and Hilary Silver. Race and job dismissals in a federal bureaucracy. *American Sociological Review*, pp. 651–660, 1992.

María J. Álvarez Peláez and Antonia Díaz. Minimum consumption and transitional dynamics in wealth distribution. *Journal of Monetary Economics*, 52(3):633–667, 2005. ISSN 0304-3932. doi: <https://doi.org/10.1016/j.jmoneco.2004.06.004>. URL <https://www.sciencedirect.com/science/article/pii/S0304393205000164>.

A More Experimental Details

This appendix section presents the robustness findings along with additional details on the simulation parameters used.

A.1 Robustness

To check for robustness, we run the same original experiment mentioned in §5.2, as follows:

The experiment is conducted using the original set of parameters in §5.1, ensuring consistency across all trials. To account for variability and assess robustness, the model is executed for 20 different randomly generated seeds. The resulting Figure 4 provides a visualization of the data by displaying the mean value, as well as error bars are included to depict the 95% confidence interval, offering insight into the robustness of the patterns shown before.




Figure 4: The simulation in 5.2 is carried out with the original parameters in 5.1. Results are for running the simulation for 20 different randomly generated seeds. The plot displays the mean value, and the error bars display a 95% confidence interval.

Analysis. The insights from Figure 4 are similar to the ones in §5.2:

- The results and consumption patterns observed before are observed in this plot, which indicates robustness.
- Workers with limited foresight and minimal knowledge of future schedule instability experience notably lower financial utility than those with greater lookahead.
- Overall, lookahead has a positive effect: workers with more foresight can better navigate financial decisions, leading to higher utility.
- Full knowledge of work schedules in advance is not necessary. Even beyond the midpoint threshold (lookahead 17), workers achieve a utility level comparable to those with complete schedule information.
- Higher income directly translates to greater utility, as individuals can consume without financial constraints. This is clearly reflected in the consistent upward shift of the plots along the y -axis.

A.2 Clarification on IQR and Other Simulation Parameters Used in §5.1

Why IQR and IQR Setting. To ensure a representative and analytically robust income distribution, we remove extreme low and high income outliers using the Interquartile Range (IQR) method. Income data is characteristically right-skewed (Benhabib & Bisin, 2018), with a small number of individuals earning disproportionately high incomes. The IQR method is particularly well-suited for such data because it is non-parametric and does not assume a normal distribution, unlike z -score-based methods, which can be distorted by skewness (Insights, 2025; Hubert & Van der Veecken, 2008; Whaley III, 2005). Specifically, per IQR’s definition, the outliers are values below $(Q1 - 1.5 \times \text{IQR})$ or above $(Q3 + 1.5 \times \text{IQR})$ where $\text{IQR} = Q3 - Q1$, effectively trimming only the extreme tails while preserving the integrity of the central distribution.

Shocks. The parameters chosen for modeling income shocks are grounded in empirical findings from real-world data. Since the shocks are very different household by household and occupation by occupation, the common practice in economics is to adopt a reasonable range.

According to the JPMorgan Chase Institute, income volatility remained fairly stable between 2013 and 2018, with households at the median experiencing a 36% change in income from month to month on average (JPMorgan Chase Institute, 2016). However, volatility levels vary substantially by household and occupation. The same report notes that income variability is highly heterogeneous, both across different families and over time for the same family. The median standard deviation of volatility is reported as 0.37.

Further illustrating this variability, retail workers often face unstable schedules that result in up to a 50% variation in work hours (and the corresponding earnings) during certain months (CLASP, 2022). Similarly, gig workers such as rideshare and delivery drivers experience changes in earnings tied to fluctuating hours and platform demand, with income shifts ranging from 3.4% to roughly 14% month-over-month depending on market conditions and app algorithms (Insider, 2024).

Taking these observations into account and to have a reasonable shock profile, we opt for income shocks that ranges from -0.4 to 0.4 , with shocks occurring as $(1 + r) \times \text{income}$, where r represents the shock value. The shocks are generated from a Bernoulli process, with the shock size parameter r uniformly sampled from $[-0.4, 0.4]$.

Discounting Factor β . Similarly, the common practice in economics is to adopt a standard, reasonable value for the discount factor, typically set at 0.95. This practice is widespread in the literature and has been employed in numerous studies (Patnaik et al., 2022; Cooper & Willis, 2014)³. However, a body of research seeks to go beyond this uniform assumption by estimating more context-sensitive β values that vary across demographic groups. For example, as discussed in the following Sections B.1, B.2, and B.3, β can change with education level or employment status. We provide a detailed analysis of these more granular real-world scenarios in Appendix B.

Return Rates. The return rate ranges used in §5.1 are grounded in empirical observations of *typical* historical asset performance of 5 years preceding 2020 across a broad spectrum (Carlson, 2021; Bank, 2020). As with most economic parameters, return rates can vary significantly depending on individual circumstances and the type of asset involved. We consider typical asset returns, not highly-volatile assets. For instance, if an individual owns cryptocurrency and other very high-risk assets, they can have return rates like +215.07% or +539.96% during certain quarters and then experience very drastic drops in asset values⁴.

Therefore, like other parameters, the common practice is to choose a range that reflects realistic variability. However, as further examined in §B.4, a more granular analysis is possible by focusing on specific typical asset classes (such as stocks or cash), which allows for a more precise assessment of return dynamics. Still, these specific return rates generally remain within the broader range defined in the initial simulation.

The description of a range and the added variance of to returns implies return rates uniformly distributed in a given range. With probability 0.1, we get an additive shock to the return rates where the additive amount

³<https://behaviouraleconomics.jasoncollins.blog/intertemporal-choice/present-bias-examples>

⁴<https://www.coinglass.com/today>

is sampled from a different uniform distribution (Thus, ‘*a range and the added variance*’ is not about uniform + Normal additive returns).

Discretization of the Dynamic Programming. In the simulations, the asset space is discretized into integral buckets with the gap between the buckets set to 1 where the number of buckets depends on the difference between highest attainable assets and the lowest possible assets (which could be non-zero depending on the income distribution).

Clarification on the Income. Income is not a constant. It is sampled uniformly from a given distribution. Agents additionally get shocks on top of income. The shock probability is set to 0.4.

Clarification on the L0+Money Intervention Setting. This setting is about adding a fixed constant to income when there is a shock. In these experiments, we did not explore multiple values, as the goal is to study the effects of giving a realistic fixed amount of money.

Details of the Simulation Runs. The main results presented in §5.2 are based on a single set of simulation runs for representative individuals within each income group and for 50 runs in the mitigation experiments. This choice was intentional, as our goal was to avoid averaging across agents, which can obscure meaningful behavioral variation, especially in a setting where individual-level dynamics are of interest. As discussed in the main text, the qualitative trends in consumption and utility remain consistent across different random seeds and repeated runs. *However, to meaningfully support this, we include robustness additional results across 20 random seeds in Figure 4, which confirm the stability of the main patterns.*

In this context, the empirical findings serve as a complement to the theoretical results established earlier, helping to illustrate broader behavioral patterns rather than focusing on fine-grained empirical fluctuations. Given the heterogeneity of agents in real-world socioeconomic environments, as shown in the detailed parameter justifications above, our interest lies in capturing these general dynamics.

B More Real-world Problems

In this appendix, we explore additional scenarios informed by real-world data and observed phenomena.

B.1 Patience Factor based on β

In models of intertemporal choice, the parameter β denotes the *subjective discount factor*, capturing the extent to which individuals prefer present consumption over future consumption. A β value between 0 and 1 indicates that future consumption is valued less than immediate consumption, with lower values reflecting a stronger preference for the present. This formulation is closely aligned with the concept of *present bias*, which describes the tendency to overvalue immediate rewards relative to delayed ones.

A lower β implies greater impatience, signifying that individuals are more inclined to choose smaller, immediate rewards over larger, delayed alternatives. In contrast, higher values of β are associated with patience and a stronger orientation toward future utility.

Empirical studies on real-world events have documented that there are economic agents who exhibit more *present-focused preferences*, meaning they are more prone to choose actions that yield immediate utility when those actions affect the present, compared to when all consequences are shifted into the future (Ericson & Laibson, 2019; Lockwood, 2020). In essence, these individuals make more impulsive decisions in the short term than they would if facing the same set of outcomes at a future point in time.

The value of β plays a critical role in determining financial behaviors such as saving and borrowing. More impatient individuals (with lower β) are typically more susceptible to accumulating debt, while those with higher β are more likely to delay gratification and engage in future-oriented behaviors like saving. In the following examples based on real-world problems, we explore how these dynamics unfold in different contexts.

B.2 Unemployment

This section examines one of the most prevalent real-world problems: job insecurity and unemployment. We focus on how unemployed individuals experience varying levels of utility based on their capacity for *lookahead*, that is, their ability to anticipate future events. Specifically, the analysis compares unemployed agents who have minimal or no information about the future against those with greater foresight (and the related β).

Although this paper primarily addresses dynamic work scheduling, the underlying framework also applies directly to unemployed individuals, who typically face profound uncertainty. An unemployed person may not know: When or whether they will find a job, whether any job obtained will be full-time or part-time, how long unemployment insurance (UI) benefits will last, whether they will experience unanticipated financial shocks during the job search, how partial employment affects UI eligibility and payments, and how and when to report earnings to maintain compliance with UI regulations.

Unemployment Insurance (UI) provides temporary income to eligible workers who lose their jobs through no fault of their own. However, *UI exhaustion* occurs once an individual has received all the benefits for which they are eligible in a given benefit year. At that point, no further payments are made unless the claim is renewed or the individual qualifies for extended benefits (Institution, 2020; on Budget & Priorities, 2021).

The actual duration of UI varies significantly by individual and state, influenced by the following factors: UI is typically calculated as a percentage of previous wages and individuals with higher earnings. Also, most U.S. states allow up to 26 weeks of standard UI benefits. However, during economic downturns or recessions, extended benefits may become available (on Budget & Priorities, 2021).

Importantly, UI benefits generally cease when an individual finds full-time employment. For part-time employment, the rules are more nuanced. If the individual remains underemployed and continues to meet the eligibility requirements (typically related to earnings thresholds), they may still receive partial benefits. However, there is uncertainty on what can be categorized as an eligible partial employment and even if they are still eligible, UI payments may be proportionally reduced based on the income earned (Inc., 2021).

This complexity creates an environment where individuals make critical consumption and job search decisions under uncertainty. Prior work by (Lockwood, 2020; Ganong & Noel, 2019) provides empirical evidence on consumption behavior during unemployment, revealing that the majority (approximately 75%) of unemployed households exhibit a discount factor $\beta = 0.9$ which indicates more patience.

Thus, to explore this example, for our simulations, we use $\beta = 0.9$, following their setup. All other simulation parameters are similar to §5.1.




Figure 5: The final utility gained for different levels of lookahead is illustrated for the four classes, each comprising 27 unemployed individuals. Agents are of similar features, with variations solely based on the temporal aspect, i.e., the amount of lookahead in their future schedules. The x -axis depicts the lookahead value and the y -axis represents the total utility at the end of T steps. Since agents are unemployed, we set β to be 0.9 based on the prior research (Lockwood, 2020; Ganong & Noel, 2019).

Analysis Figure 5 illustrates the impact of foresight on utility for unemployed individuals across income levels. Three main insights emerge:

1. Individuals with little to no information about the future face significant utility losses. This is expected, as uncertainty in income, benefit eligibility, and job prospects limits their ability to plan and allocate resources effectively.
2. With increased foresight, individuals make better-informed financial decisions. The utility gain with more lookahead demonstrates that even in uncertain environments, information helps optimize consumption and savings strategies.
3. Higher-earning individuals consistently achieve greater utility, which aligns with their ability to consume without constraint. This result is evident from the upward shift in utility across the y -axis for the higher-earning group.


B.3 Real-world Educational-based Patience Factors

In this section, we examine another real-world scenario: the relationship between educational attainment, consumption behavior, and the benefits of having lookahead. While the previously used discount factor of $\beta = 0.95$ in §5.1 represents a standard and commonly accepted value, our earlier analysis on unemployment illustrates that patience can vary across different demographic groups. This motivates a more context-sensitive exploration of β , particularly as it relates to educational background.

Prior empirical research demonstrates that individuals’ time preferences reflected in their β values are systematically associated with their level of education. Specifically, individuals with a high school diploma but no college degree tend to exhibit more impatient behavior, corresponding to a lower β of approximately 0.50.

Conversely, those with a college degree tend to be more patient, with an estimated β of 0.83 (Laibson et al., 2024; Lockwood, 2020). Therefore, in this simulation, we adopt $\beta = 0.50$ for workers with mid-level education and $\beta = 0.83$ for those with higher education, to better align our model with observed real-world behavior.

Except for this variation in β values based on educational attainment, the experimental setup remains consistent with the baseline configuration described in §5.1. Figure 6 presents the resulting utility outcomes for workers of varying education levels.



(a) High school degree lookahead analysis

(b) College degree lookahead analysis

Figure 6: The final utility gained for different levels of lookahead is illustrated for four income classes, each comprising 27 individuals. Workers are of similar features, with variations solely based on the temporal aspect, i.e., the amount of lookahead in their work schedules. The x -axis depicts the lookahead value, and the y -axis represents the total utility at the end of T steps. Here, in the left plot, the income classes have a high school education level with $\beta = 0.5$ and they have higher education on the right plot with $\beta = 0.83$.

Analysis. The insights derived from Figure 6 can be summarized through several key observations concerning workers with lower educational attainment:


- Workers with lower education levels tend to exhibit higher levels of impatience. As a result, they reach their peak attainable utility more quickly than their more patient counterparts. For instance, individuals with only a high school education converge to their maximum utility within approximately 5 weeks, whereas college-educated workers continue to gain from lookahead information for up to 10 weeks.
- As observed in earlier settings, limited foresight leads to significantly lower utility. Workers who lack access to future information or are unable to anticipate potential instability in their schedules are less able to optimize their consumption, resulting in diminished financial well-being.
- Greater lookahead consistently improves utility outcomes. The ability to foresee and plan around future events allows workers to adjust their financial decisions accordingly (particularly within the first 5 weeks for high school-educated workers and the first 10 weeks for those with college degrees).
- Higher income levels and higher education predictably lead to higher utility even without lookahead. This is reflected in the difference in the range of the utility values along the y -axis in the two plots.

B.4 Real-world Return Rates on Liquid Assets


Similar to §5.3, this section’s goal is to investigate the impact of asset appreciation and depreciation on decision-making across various levels of lookahead, i.e., some workers are already at a (dis)advantage in terms of assets. While the return rates discussed in §5.3 are accurate when considering a broad spectrum of assets and realistic return ranges (Carlson, 2021; Bank, 2020), we can refine this analysis by focusing on specific categories, particularly liquid assets. Liquid assets refer to holdings that can be readily converted into cash with minimal loss in value.

In this section, we examine two of the most commonly held liquid assets: *cash* and *stocks*. The primary distinction between them lies in their nature: stocks are investment vehicles that yield returns over time, whereas cash represents direct monetary holdings. In contrast, cash does not appreciate and is vulnerable to inflationary erosion but represents the most liquid form of asset ownership ⁵.

To model realistic return variability, we use empirical data from the five years preceding 2020 (i.e., 2016–2020) (Carlson, 2021; Bank, 2020). For stocks, we set a return range of 10%. For cash-only holdings, we assume the individual is holding cash only, so based on real-world returns (Carlson, 2021; Bank, 2020), we use a tighter range of -0.5% . All other experimental conditions are held similar to 5.3.



(a) Stocks return rate analysis



(b) Cash return rate analysis

Figure 7: Individuals with similar features but varying levels of lookahead under various return rates. The return rates are set to 1.1 for stocks in the left plot. In the right plot, the return rates are set to 0.995 for cash. Other simulation parameters remain the same as 5.1.

⁵<https://www.investopedia.com/articles/investing/103015/cash-vs-bonds-what-pick-times-uncertainty>

Analysis. The results, depicted in Figure 7, are consistent with those found in §5.3 and offer several noteworthy insights:

- Individuals achieve higher utility under favorable return conditions, like stocks. This trend underscores the intuitive relationship between asset performance and consumption outcomes.
- When holding a *very mildly* depreciating asset such as cash (only -0.5% depreciation), lookahead still boosts utility. Individuals quickly adapt their consumption in anticipation of diminishing purchasing power. Notably, since this asset is depreciating very mildly, the individuals attain near-maximum utility later than the original negative return setting in §5.3 and by lookahead 15.
- In scenarios with anticipated positive returns (e.g., stocks), the knowledge of a 10% increase helps individuals a lot in the beginning. However, after lookahead 5, individuals can afford to delay consumption, knowing they are likely to benefit from capital appreciation. As a result, convergence to peak utility occurs more gradually across all income groups (very late and around lookahead 20 here), reflecting increased consumption flexibility.