

Reinforcement Learning Paycheck Optimization for Multivariate Financial Goals

Melda Alaluf Giulia Crippa* Sinong Geng[§] Zijian Jing Nikhil Krishnan*
Sanjeev Kulkarni^{*,‡} Wyatt Navarro Ronnie Sircar* Jonathan Tang

April 30, 2023

Abstract

We study paycheck optimization, which examines how to allocate income in order to achieve several competing financial goals. For paycheck optimization, a quantitative methodology is missing, due to a lack of a suitable problem formulation. To deal with this issue, we formulate the problem as a utility maximization problem. The proposed formulation is able to (i) unify different financial goals; (ii) incorporate user preferences regarding the goals; (iii) handle stochastic interest rates. The proposed formulation also facilitates an end-to-end reinforcement learning solution, which is implemented on a variety of problem settings.

1 Introduction

We propose a reinforcement learning solution to paycheck optimization. Specifically, one aims to allocate monthly income in order to achieve goals like paying out loans, purchasing a mortgage, saving for retirement, etc. Indeed, such a problem is common in everyday life and similar services are provided by various companies. In this work, we hope to provide a rigorous framework for such problems with a reinforcement learning solution.

Finding a suitable problem formulation for paycheck optimization is especially challenging. First, the goals of paycheck optimization are often multivariate and quite heterogeneous. Therefore, it is cumbersome to unify such goals and optimize them simultaneously. Second, incorporating the preferences of users becomes especially complicated for paycheck optimization. For instance, some goals like paying out credit card debt can be more urgent to the user than saving money or purchasing a mortgage, while the amounts vary greatly. It is unclear how to incorporate such information into decision-making methods. Third, the interest rates of financial goals (e.g. inflation rate, savings rates, etc.) evolve stochastically over time. This makes learning an optimal paycheck allocation strategy even more challenging. Finally, without a proper formulation, the powerful decision-making tools in machine learning and control are not applicable to this problem.

To the best of our knowledge, a quantitative solution for paycheck optimization is missing. Existing results on paycheck optimization are mainly analytical without an implementable methodology (Swart, 2004; Archuleta & Grable, 2011; Hershey et al., 2013). One can also consider paycheck optimization as a non-traditional robo advising problems with different targets. However, existing

*ORFE Department, Princeton University, Princeton, NJ 08544, USA

[§]Computer Science Department, Princeton University, Princeton, NJ 08544, USA

[‡]Electrical & Computer Engineering Department, Princeton University, Princeton, NJ 08544, USA

literature is mainly for portfolio optimization (Giudici et al., 2022; D’Acunto et al., 2019; D’Acunto & Rossi, 2021) or other single financial goals (D’Acunto & Rossi, 2021): such methods are not applicable to paycheck optimization with the multiple heterogeneous goals studied in this work. Some existing paycheck optimization solutions rely on a simple waterfall method. Specifically, the user needs to prioritize different goals in an absolute order to finish the goals one by one. In other words, all incomes will be allocated to a specific goal and only when one is met will the next one be considered. As a result, the method is incapable of targeting multiple goals simultaneously and thus is generally sub-optimal. An example where the waterfall method performs poorly is given in Appendix A.

Separately, there exists a huge amount of literature on reinforcement learning for decision-making in various scenarios. Such methods provide flexible solutions for many different decision-making problems, but are not directly applicable to paycheck optimization.

In this work, we propose a utility maximization framework for paycheck optimization and a data-driven policy gradient method. First of all, we formulate the paycheck optimization as a utility maximization problem to unify various financial goals and incorporate user preferences. In detail, we leverage piecewise-linear utility functions. Whenever a goal is active (i.e. it is still beneficial to allocate income to this goal) the corresponding utility function is negative, while it becomes zero otherwise. This design has two advantages: on one hand, it encourages a policy to finish the goals; on the other hand, it is possible to express the user-specific preference for each goal via the slope of the utility function - the steeper the slope, the more beneficial it is to allocate income to the corresponding goal.

The decision-making target is to maximize the sum of the utility functions of each goal across time. With the proposed utility maximization framework, we conduct policy gradient to solve for an optimal paycheck allocation strategy. Specifically, with the collected data, we implement policy learning using gradients estimated from the data. As a result, we learn a paycheck optimization policy in a data-driven and model-free manner, without specifying any stochastic model. This provides a flexible solution to paycheck optimization.

2 Problem Formulation

Paycheck optimization studies the problem of income allocation over different financial goals. Examples of financial goals include paying out credit card debt, paying out student loans, saving for a home down payment, saving emergency funds, or saving for retirement using for example 401Ks or IRAs. At time t , we use S_t to denote a user’s income, and π_t to denote the fraction of S_t assigned to different financial goals. By optimizing the income allocation π_t , we aim to complete all the financial goals.

In paycheck optimization, the financial goals are heterogeneous. For instance, savings goals like retirement and emergency funds are different from debt goals like student loan and credit card debt. Indeed, the former depend on interest rates that increase the value of the wealth assigned to them, i.e. they contribute to finishing them. On the other hand, the latter’s interest rates increase the value of the debt itself, delaying their completion. Goals also differ based on their maturity. Short-term goals, like credit card debt, have much higher interest rates and thus, should be treated more urgently with respect to longer-term goals. In practice, it is unclear how to unify such heterogeneous traits and optimize these goals simultaneously.

While the main objective of paycheck optimization is finishing all financial goals, it is also important to consider the users’ preferences. Instead of completing the goals as fast as possible, different users may have different priorities for each goal. For example, some users could prefer

saving for purchasing a house more than saving for retirement. On the other hand, some might prioritize retirement, given that such investments may have higher interest rates that would allow them to finish all other goals more quickly. Additionally, an agent may choose to pay down high interest debt first, or they may prioritize the confidence boost that comes from zeroing out debt and choose to pay out small debts first. As a result, how to quantify the preference of users and incorporate them into paycheck optimization is also an open question.

3 Paycheck Optimization as Utility Maximization

In this section, we formulate paycheck optimization as a utility maximization problem. For each financial goal, we define (i) a state variable, (ii) its dynamics and, (iii) a utility function.

3.1 State variable

To devise a utility maximization objective, we first define the state variable for this problem. Let I be the set of financial goals with cumulative totals that we aim to achieve. Then, $X_t \in \mathbb{R}^{|I|}$ is the state variable, i.e., the fraction of each goal that still needs to be paid out at time point $t \in \{0, \dots, T\}$. For $i \in I$, the component X_t^i gives a normalized measure of the proportion of goal i that is left to complete, where $X_t^i = 0$ denotes that the goal has been completed by time t , and $X_0^i = 1 \forall i \in I$ since we start off without having contributed income towards any of the goals.

3.2 Dynamics

Critically, the specific dynamics of X_t are different for each financial goal. We use S_t and π_t^i to denote the monthly income (increasing at the rate of inflation) and the fraction of income assigned to goal i at time t respectively. For goals involving debt repayment like credit card debt and student loans, the dynamics of X_t^i is given by

$$X_{t+1}^i = (1 + r_t^i)X_t^i - \frac{S_t \pi_t^i}{G^i}, \text{ with } i \in \{\text{Credit Card Debt, Student Loans}\} \quad (1)$$

where r_t^i is the interest rate for goal i and G^i is the total amount for goal i . In words, X_t^i is updated according to the interest rate, before subtracting the proportion of the goal that the user will pay at time t . In practice, the interest rates and goal amount are different for each financial goals (see Table 2).

For financial goals involving savings like home down payment and emergency funds, we define

$$X_{t+1}^i = 1 - (1 + r_t^i)(1 - X_t^i) - \frac{S_t \pi_t^i}{G^i}, \text{ with } i \in \{\text{Home Down Payment, Emergency Funds}\} \quad (2)$$

which differs from the previous case, since the interest rate contributes to completing the different goals. In general, the emergency fund will not have an interest rate, since it is intended to be easy and accessible cash on-hand. However, we assume that mortgage savings are invested in a risk-free financial instrument with some constant rate of return, while in Section 6 we will consider stochastic rates for this investment.

Finally, for retirements savings, the structure is considerably different due to the presence of tax-advantaged savings accounts which are commonly used for retirement, i.e., 401K and IRA. If goal i is the retirement savings, denoted as RS , then

$$X_{t+1}^{RS} = 1 - (1 + r_t^{RS})(1 - X_t^{RS}) - \frac{m_t + S_t^{RS}(\pi_t^{401K} + \pi_t^{IRA} + \pi_t^{RS})}{G^{RS}}$$

where r^{RS} is the rate of returns on retirement savings (which can again be extended to evolve stochastically). Note that $(1 - X_t^{\text{RS}})$ is the fraction of retirement savings which has been already paid out, and so $1 - (1 + r_t^{\text{RS}})(1 - X_t^{\text{RS}})$ gives the fraction of debt still outstanding after applying the rate of return. Eventually, we subtract the amount paid to the retirement goal, as well as the amount paid to 401K and IRA (since they also contribute to retirement). Also, we deduct a factor m which represents employer matching for 401K, which depends on the level defined in the employer’s specific plan.

3.3 Utility

For each goal we wish to define a corresponding utility function $u_i : \mathbb{R} \rightarrow \mathbb{R}$. We define the utility function as a function of the fraction finished in each financial goal, denoted as $\bar{X}_t^i = 1 - X_t^i$. To formulate the priority of task i , specified by users, we use two positive constants p^i and q^i , where a larger p^i (or q^i) corresponds to a greater urgency to complete goal i . Specifically, with $\bar{X}_t^i = \bar{x}$, $p^i = p$, $h^i = h$, and $q^i = q$, we define two types of utility functions:

$$w_1(x; p) = -p \cdot \max(0, 1 - \bar{x})$$

$$w_2(x; p, q, h) = -q \cdot \max(0, 1 - \bar{x}) - (p - q) \cdot \max(0, 1 - \bar{x} - h),$$

where w_1 represents the utility for a single-phase goal, while w_2 a two-phase one. This allows us to specify multiple priorities p and q , depending on the proportion of goal already been met. Note that w_2 includes a parameter $h \in [0, 1]$ to specify the crossover point between when a user pays off the first segment of the goal and moves onto the second one. Below we give a plot of such utility functions.




Figure 1: Different utility functions

Let us now note the following about the defined utility functions. First, both are continuous with finite derivative, which allows for gradient-based methods to solve the corresponding maximization problem (3) below. Second, they return negative values when $x > 0$, but stay at zero when $x \leq 0$. Moreover, as stated above, the value of p, q for different financial goals model the user’s preference with respect to each goal. Specifically, p, q are non-negative weights which can be chosen to incentivize completion of goals: the larger the value of p , the more incentive to complete the first (or only) segment of the goal; the larger the value of q , the greater the incentive to complete the second segment.

For the single phase goals, e.g. $i \in \{\text{Credit Card Debt, Student Loans, Home Down Payment, Retirement Savings}\}$, we have

$$u_i(X_t^i) = w_1(X_t^i; p^i),$$

where p^i is the weight for goal i . For the emergency fund we have

$$u_{\text{EF}}(X_t^{\text{EF}}) = w_2\left(X_t^{\text{EF}}; p^{\text{EF}}, q^{\text{EF}}, h^{\text{EF}}\right),$$

where $(1 - h^{\text{EF}})G^{\text{EF}}$ and $h^{\text{EF}}G^{\text{EF}}$ are the amounts to be completed with urgency p^{EF} and q^{EF} respectively.

In addition to the goals with state variables, we also want to assign utilities to the 401K and IRA goals, since making regular contributions to these accounts could be a user desired goal in addition to their contribution to the retirement savings goal. Since there are no cumulative totals to meet, the utility for the 401K and IRA will be assessed based on the contribution at each time t .

401K is a two-phase goal, where the first phase consists of contributions up to a minimum level, while the second one comprises contributions up to the maximum allowed level. Hence, for 401K the utility will be given by

$$u_{401K}(X_t^{401k}) = w_2\left(1 - \frac{\pi_t^{401K}}{M^+}; p^{401K}, q^{401K}, \frac{M^+ - M^-}{M^+}\right),$$

where M^+, M^- are the maximum and minimum income percentage contribution levels to 401K respectively. For IRA instead,

$$u_{\text{IRA}}(X_t^{\text{IRA}}) = w_1\left(1 - \frac{S_t \pi_t^{\text{IRA}}}{I^+}; p^{\text{IRA}}\right),$$

where I^+ is the maximum permitted income contribution to the IRA.

Eventually, with utilities defined for every and each goal, we can define the paycheck optimization target as maximizing the expected total utility:

$$\max_{\pi_t} V(\pi_t) \quad \text{with} \quad V(\pi_t) = \sum_{t=0}^T \sum_{i \in I} u_i(X_t^{i,\pi}), \quad (3)$$

where the evolution is over the dynamics of $X_t^{i,\pi}$, under policy π_t .

Note that, while the proposed framework provides the ability to assign a different preference (weights) to every goal, it is always feasible to fix some of them to simplify the problem. This is vital from a practical perspective, since it might be difficult for users to order the importance of all financial goals. For instance, we might assume that the user have the same preference for all debt goals and thus fix p^i for student loans and credit card debt as one value. In practice, this can simplify the communication with users when trying to come up with p^i 's and q^i 's.

4 Deep Deterministic Policy Gradient

We aim to solve (3) by deep deterministic policy gradient (Silver et al., 2014). Specifically, we parametrize the policy as a deep neural network of X_t :

$$\pi(t) = f(X_t; \theta), \quad (4)$$

where f denotes the neural network with parameter θ . Then, for each iteration, we update the parameter θ by the estimated gradients of $V(\pi_t)$ with respect to θ .

Our procedure is as follows:

Algorithm 1 Paycheck optimizer

- 1: **Input:** Amounts for each goal, interest rates, matching percents, number of iteration N .
 - 2: **for** $n \in [N]$ **do**
 - 3: Initialize X_0 as one for each goal.
 - 4: For each month t :
 - (a) $\pi_t = f(X_t; \theta)$
 - (b) Update X_t from X_{t-1} according to the rules of Section 3.2
 - 5: Estimate $\frac{dV(\pi_t)}{d\theta}$
 - 6: Maximize $V(\pi_t)$ with respect to θ
 - 7: **end for**
 - 8: **Output:** Return θ
-

Where Step 1 is initialization for each training epoch, Step 4(b) follows (4), and Steps 5,6 are where we optimize our neural network. Optimization can be done with any standard algorithm, for example gradient descent or ADAM (Kingma & Ba, 2014).

5 Simulation with Constant Rates

In this section, we implement the proposed method for paycheck optimization. We aim to show that our method is readily available for paycheck optimization for users with different preferences. In the following, we first describe the experiment protocol and then provide results.

Protocol To implement the proposed method, we consider three types of users with different preferences for each of the financial goals:

- the home buyer, whose priority is purchasing a home as quickly as possible;
- the saver/retirement planner, whose priority is maximizing retirement savings and saving for emergency;
- the debtor, who prefers to pay off debt first.

For each category, we construct a representative user, with preference weights selected reflecting the user type (see Table 1). For each of the users, we set the same input data (see Table 2). With this experiment, we demonstrate that the proposed paycheck optimization framework is able to effectively address the preferences of users, while finishing each financial goal in an efficient manner.

Results For each representative user, we report the contribution to each goal over time under the learned policy in Figure 2. We note that each goal is successfully completed under the learned paycheck allocation policy. Also, the result is consistent with the user preference. Specifically, home buyer does buy a mortgage earlier than others; the savings of saver grows faster; debtor pays off the debts in the fastest manner.

Users	Preference Weights
Home Buyer	Set $p^{\text{mortgage}} = 20.0$ and other preference weights as one.
Saver	Set $p^{\text{retirements}} = p^{401k} = q^{401k} = p^{IRA} = 20.0$, $p^{\text{emergency fund}} = 5.0$, $q^{\text{emergency fund}} = 3.0$, and other preference weights as one.
Debtor	Set $p^{\text{credit card}} = p^{\text{student loan}} = 20.0$, $p^{\text{emergency fund}} = 5.0$, $q^{\text{emergency fund}} = 3.0$, and other preference weights as one.

Table 1: Preference weights for different users.

Input parameter	Value
Monthly income	\$7,500
Inflation rate	2%
Stock Market Rate of Return	10%
Credit Card Debt	\$825
Credit Card APR	20%
Student Loan Debt	\$80,000
Student Loan APR	4%
Mortgage Down Payment	\$157,000
Emergency Amount (I)	\$1,800
Emergency Amount (II)	\$9,000
Retirement Savings	\$1,000,000
IRA (monthly contribution)	\$500
401K (min and max contribution)	6% – 13% of salary
Time-horizon	10 years

Table 2: Inputs for each user.

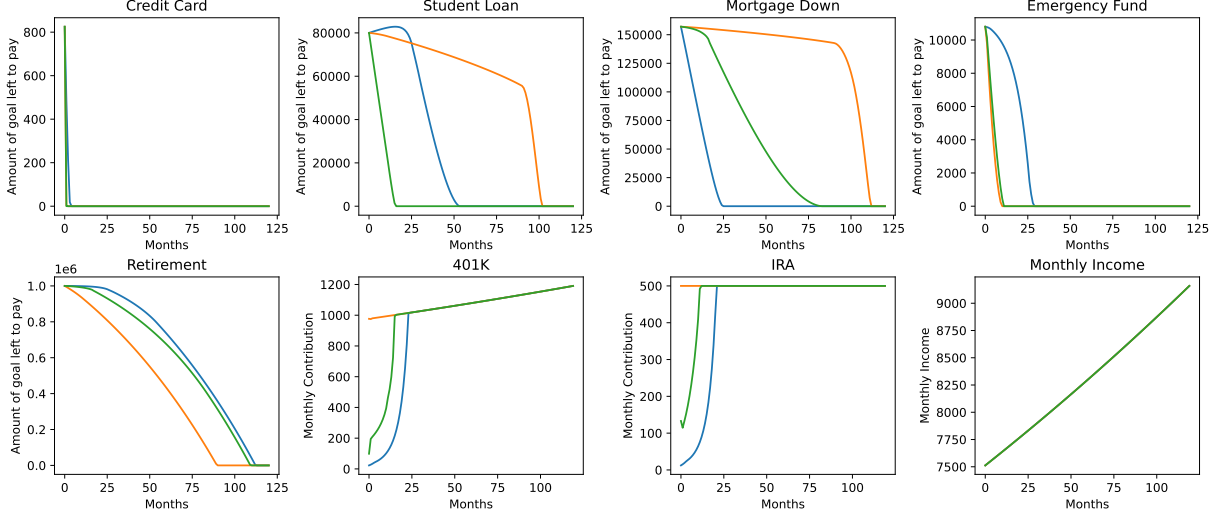


Figure 2: Contribution to each goal over time under the learned policy with constant rates for three different representative users: home buyer in blue, saver in orange, and debtor in green

Explainability From the experiments above, we can explain the learned policy by preference weights. Specifically, the policy function prefers to finish the financial goals with higher preference weights by allocating more income to such goals. One can further examine the effects of each state variable on the learned policy by using Shapley values. Some examples in portfolio optimization include Babaei et al. (2022); Colini-Baldeschi et al. (2018). However, it is nontrivial to extend such analysis from portfolio optimization to multiple financial goals in our setting. We thus defer that to future work.

6 Extension to Stochastic Rates

In the previous analysis, we fixed the rates $\{r_t^i\}_{i \in I}$ as constant over time, while in practice some or all the rates may evolve stochastically over time. In this section, we extend our method to the stochastic rate case. Specifically, we assume that the rates $\{r_t^i\}_{i \in I}$ follow a Markov process, so that the dynamics of X_t^i are also stochastic. As a result, the value function in (3) needs to be redefined as

$$V(\pi) = \mathbb{E}_\pi \left[\sum_{t=0}^T \sum_{i \in I} u_i(X_t^i) \right], \quad (5)$$

where the expectation is over the dynamics of X_t^i dependent on r_t^i , under the policy π .

We maximize the value function (5) following the deep deterministic policy gradient in Section 4 while using data of rates. We use $\tau = \{r_t^i\}_{t=0, i \in I}^T$ to denote a data trajectory of rates. Let $\{\tau_k\}_{k=1}^n$ denote a dataset with n observed trajectories. We parameterize the policy function as deep neural network with parameter θ :

$$\pi(t) = f(x_t, \{r_t^i\}_{i \in I}; \theta). \quad (6)$$

Then, we train the neural network f by maximizing the sample-average utility function

$$\frac{1}{n} \sum_{k=1}^n \sum_{t=0}^T \sum_{i \in I} u_i(x_t^{i, \theta, k}), \quad (7)$$

where $x_t^{i,\theta,k}$ denotes the state value at time point t under the policy function (6) with stochastic rates following τ_k . Thus, the gradient of (7), with respect to θ , is derived as

$$\frac{1}{n} \sum_{k=1}^n \sum_{t=0}^T \sum_{i \in I} \frac{du_i(x_t^{i,\theta,k})}{d\theta}.$$

In other words, in each iteration of our policy learning, we use the average over n trajectories to calculate the gradient and update θ . Note that our procedure does not need independence assumptions or parametric models for the stochastic rates $\{r_t^i\}_{i \in I}$: the dynamics of $\{r_t^i\}_{i \in I}$ is purely handled by data.

7 Simulation with Stochastic Rates

In this section, we conduct experiments for the case with stochastic rates. Under the same setup as the experiments in Section 5, we treat rates r_t^i as stochastic processes instead of constants. Following the procedure in Section 6, to handle gradient estimation for (5), we select the Consumer Price Index (CPI), the market yield of 3-Month U.S. Treasury Bills and the S&P500 Index return from 1985-2022 in order to model inflation, as well as the rate of return on mortgage down payment savings and retirement savings respectively. Under the learned policy, the contributions over time for different users are reported in Figure 3 over a ten year time horizon from 2012-2022. Note that

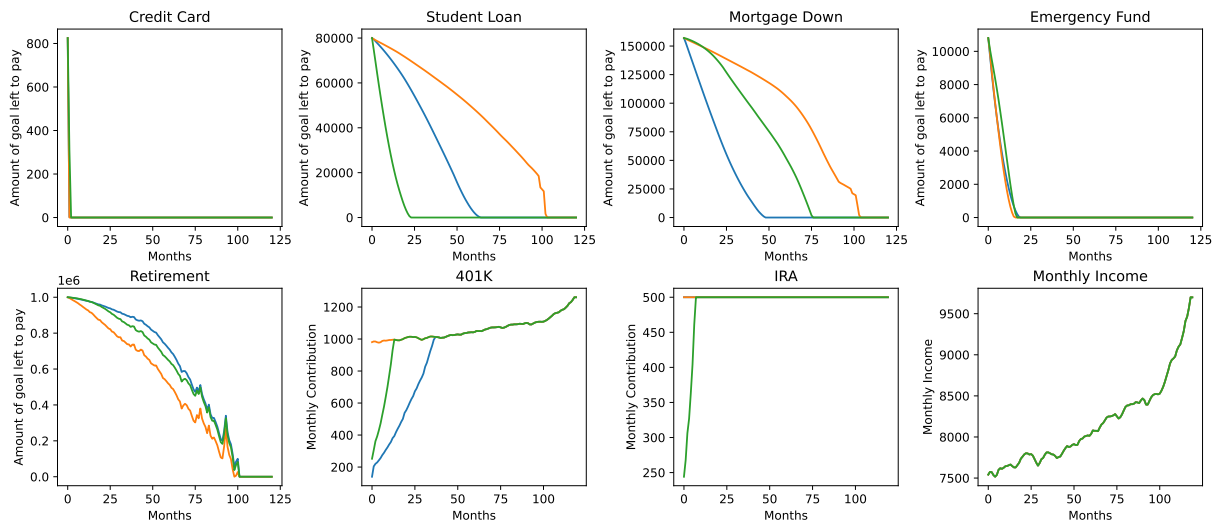


Figure 3: Contribution to each goal over time under the learned policy with stochastic rates for three different representative users: home buyer in blue, saver in orange, and debtor in green. Note that the monthly income suffers a sharp increase after month 100, since it is directly affected by inflation (which has hiked over the last couple of years).

the results are mainly consistent with the deterministic rates case in Section 5. However, here the contributions to each goal have more fluctuations, since the policy is implicitly trying to predict the potential rate change and adjust the paycheck assignment accordingly.

8 Conclusion

We propose a framework for paycheck optimization with an end-to-end reinforcement learning solution. By formalizing the problem into a piecewise linear utility maximization problem, our method is able to handle heterogeneous financial goals, the preferences of users, and also the stochastic rates. We empirically demonstrate the applicability of the proposed method.

A Example of waterfall failure

Consider an individual with disposable income of \$1000 a month and two financial goals, each of which involve paying off debt:

- Goal 1 is to pay off \$1000 with no interest rate, and with priority $p^1 = 1000$.
- Goal 2 is to pay off $\$ \frac{1,000,000}{1+r}$ with interest rate $r = 0.001$ and priority $p^2 = 1$.

Recall that the waterfall method consists in paying off the goals in order of priority. Accordingly, the strategy would be as follows. First, we would pay off Goal 1, since $p^1 > p^2$. Hence, $X_0^1 = 1000$, $X_t^1 = 0 \forall t > 0$. Notice that, for goal 2 we have $X_0^2 = \frac{1,000,000}{1+r}$, while for goal 1, $X_1^2 = 1,000,000$. For each subsequent time $t > 1$, notice that the increase in X^2 due to interest will be equal to 1000. Hence, at each time step t , the debt for Goal 2 will increase by 1000, resulting in the user needing to allocate all her paycheck to pay down this increase. Thus, $X_t^2 = 1,000,000$ for $t > 0$, and the user will never be able to pay down Goal 2.

On the other hand, let us consider a strategy where we recognize the threat of future compounding interest. The user would optimally split her paycheck evenly among the two goals. Therefore, $X_0^1 = 1000$, $X_1^1 = 500$, $X_t^1 = 0$ for $t \geq 2$. For Goal 2, $X_0^2 = \frac{1,000,000}{1+r}$, $X_1^2 = 999,500$, $X_2^2 = 999,999.5$. From this point onward, the user would allocate all her paycheck, equal to \$1000 towards Goal 2, which will be gradually paid down. Indeed, the increase in the debt due to interest will always be smaller than her paycheck.

References

- Kristy L Archuleta and John E Grable. The future of financial planning and counseling: An introduction to financial therapy. In *Financial planning and counseling scales*, pp. 33–59. Springer, 2011.
- Golnoosh Babaei, Paolo Giudici, and Emanuela Raffinetti. Explainable artificial intelligence for crypto asset allocation. *Finance Research Letters*, 47:102941, 2022.
- Riccardo Colini-Baldeschi, Marco Scarsini, and Stefano Vaccari. Variance allocation and shapley value. *Methodology and Computing in Applied Probability*, 20:919–933, 2018.
- Francesco D’Acunto and Alberto G Rossi. New frontiers of robo-advising: Consumption, saving, debt management, and taxes. *Saving, Debt Management, and Taxes (February 1, 2021)*, 2021.
- Francesco D’Acunto and Alberto G Rossi. Robo-advising. In *The Palgrave Handbook of Technological Finance*, pp. 725–749. Springer, 2021.
- Francesco D’Acunto, Nagpuranand Prabhala, and Alberto G Rossi. The promises and pitfalls of robo-advising. *The Review of Financial Studies*, 32(5):1983–2020, 2019.

- Paolo Giudici, Gloria Polinesi, and Alessandro Spelta. Network models to improve robot advisory portfolios. *Annals of Operations Research*, pp. 1–25, 2022.
- Douglas A Hershey, Joy M Jacobs-Lawson, and James T Austin. Effective financial planning for retirement. 2013.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. URL <https://arxiv.org/abs/1412.6980>.
- David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pp. 387–395. PMLR, 2014.
- Nico Swart. *Personal financial management*. Juta and Company Ltd, 2004.