Subsidy Allocations in the Presence of Income Shocks

Rediet Abebe,[1, 2] Jon Kleinberg,[2] S. Matthew Weinberg[3]
[1]Harvard University
[2]Cornell University
[3]Princeton University

Abstract

Poverty and economic hardship are understood to be highly complex and dynamic phenomena. Due to the multi-faceted nature of economic welfare, assistance programs targeted at alleviating hardship can face challenges, as they often rely on simpler measures of welfare, such as income or wealth, that fail to capture to full complexity of families' state. Here, we explore one important dimension – *susceptibility to income shocks*. We introduce a model of welfare that incorporates income, wealth, and income shocks. We analyze this model to show that it can vary, at times substantially, from measures of welfare that only use income or wealth. We then study the algorithmic problem of optimally allocating subsidies in the presence of income shocks. We consider two well-studied objectives: the first aims to minimize the expected number of agents that fall below a given welfare threshold (a *min-sum* objective) and the second aims to minimize the likelihood that the most vulnerable agent falls below this threshold (a *min-max* objective). We present optimal and near-optimal algorithms for various general settings. We close with a discussion on future directions on allocating societal resources and the ethical implications of related approaches.

1 Introduction

Understanding and measuring economic hardship is a fundamental question that directly informs the design of policies and assistance programs aimed at addressing the needs of vulnerable individuals and families (Atkinson 2003). A crucial challenge is the range of factors that play a role in poverty and economic hardship, including health, demographics, social ties, intergenerational dynamics, and many other dimensions (Grusky 2018).

Recent studies have sought to address the gap between official measures of welfare and the more complex formulations that might be needed to accurately identify the sources of greatest need. One active and ongoing effort is the *Poverty Tracker* program (Wimer et al. 2014), which surveys approximately 2300 families in New York City and documents the intricate associations between their circumstances and levels of hardship. As with other studies in this area, it is in part based on the premise that we need to broaden our frameworks for quantifying economic well-being; the researchers involved in the program write that "official 'income only' measurements of poverty ... painted a picture that was too optimistic and didn't capture the magnitude of disadvantage, nor the true struggles New Yorkers face in trying to make ends meet" (Wimer et al. 2014).[1]

Shocks. What are the missing dimensions in these basic measurements? Several additional components of economic hardship manifest themselves through a common mechanism: unexpected and disruptive *shocks* to a family's economic state. Such shocks can be a result of an unexpected expense (e.g., a parking ticket or a health bill caused by an accident), a delayed paycheck or unexpected loss of a job, a dissolution or loss of a romantic or other close personal relationship, interactions with the criminal justice system, and many other experiences. Income shocks are receiving increasing attention from social scientists and policy-makers alike; summarizing a recent round of findings, the Poverty Tracker analysis discussed above reports that "the most persistently disadvantaged New Yorkers are beset by repeated shocks to their finances and well-being" (Wimer et al. 2014).

A crucial point is that families vary significantly in their ability to withstand income shocks; while an unexpected bill might be a mere inconvenience for some families, for other families, it can lead to eviction, poor health, loss of a job, and other undesirable outcomes that may trigger or lock families into persistent poverty (Atake 2018; Desmond 2016; O'Flaherty 2009; Shapiro and others 2004). In many cases, it is significantly more challenging to remedy the consequences of such experiences than it is to prevent families from experiencing them in the first place.

Despite the centrality of shocks in hardship, they do not play a correspondingly central role in the evaluations and decisions made by social assistance programs. For instance, standard eligibility guidelines for housing assistance programs are based on income, adjusted for family size, or percentage of income spent on housing. Likewise, other assis-

[1]The Poverty Tracker study is based on surveys of families of all income levels across New York City and is conducted by the Robin Hood Foundation and the Columbia University Population Research Center.

tance programs such as the Supplemental Nutrition Assistance Program and Medicaid are based on income eligibility. Yet two families that look similar under such measures may still differ significantly in their vulnerability if one family experiences a significantly different profile of shocks.

There is thus a danger of misprioritizing, if we are not taking into account factors that we know to be crucial. What would it look like to incorporate information about shocks into disadvantage determinations, and how might it inform decisions about assistance?

The present work. We develop a stylized model for the state of an agent (representing a family potentially in need of assistance) as they experience shocks over time. We then use this model to formulate the problem of allocating subsidies to agents, when the total amount of subsidy is constrained by a given budget. One option for approaching this problem is to take into account only the income of each agent. But, given that shocks can have a significant effect on families' welfare, how should we incorporate information about shocks into the search for allocations? And, how much can our allocation decisions change when we use this information about shocks?

We tackle the problem of modeling agents' welfare and using these to study allocation problems where the planner has a fixed set of of funds across the agents. We consider models in which the agents can receive funds in the form of income or wealth subsidies. We study the problem of optimally allocating such subsidies to optimize for different objective functions, and present optimal and near-optimal algorithms for a variety of natural settings. In the process, we also obtain a number of insights into the structure of the problem itself. In particular, the algorithms we obtain turn out to have a natural structure based on priority orderings on the agents, which we believe to be interesting objects in and of themselves. These insights highlight the ways in which the priority orderings used to target agents can change depending on the objective function and the subsidy type. We close with a discussion on open questions as well as societal considerations when using optimization based methods to inform assistance programs.

2 Allocating Wealth and Income Subsidies

We begin by specifying the theoretical model and problem formulation. The model is based on the structure of *ruin processes* that are standardly used to represent risk in insurance markets (Asmussen and Albrecher 2010); to make the exposition self-contained, it is useful to describe our version of the model from first principles. The optimization problems we study, based on interventions to modify the ruin probabilities, are natural given the motivation in the previous section, but less standard in the earlier literature as it is focused on insurance markets rather than poverty and economic vulnerability. We describe these problems in the latter part of the section, after first establishing the basic model.

A Model of Income, Reserve, and Shocks. There are n agents; we can think of each as representing a family that a planner would like to assist. Each agent i has an (net) income c_i per unit time, and an initial reserve u_i. Income can be thought of as the difference between the agent's earned income minus their expenses during each time period. Time runs continuously; so, in the absence of shocks, agent i's reserve after an amount of time t would be $u_i + c_i t$.

The shocks experienced by agent i operate as follows: shocks arrive at randomly selected discrete points in time $T_{i1}, T_{i2}, \ldots$ with the gaps between them $T_{i(j+1)} - T_{ij}$ distributed independently according to some distribution G_i. Thus, if a shock happens at some time T, we can imagine setting an independent random "countdown timer" of length drawn from G_i; when this timer expires, the next shock happens. When the j^{th} shock happens, it has a size S_{ij} drawn from a *shock-size distribution* F_i, and S_{ij} is subtracted from the agent's current reserve.

For concreteness, unless stated otherwise, assume that the shocks arrive according to a *Poisson process*, which has the structure described above, with the gaps between consecutive shocks drawn from a distribution G_i that is exponentially distributed with rate β_i. The expected length of the gap between consecutive shocks is $1/\beta_i$; we can equivalently think of this as saying that there are β_i shocks per unit time in expectation. The use of an exponential distribution for the gap between consecutive shocks yields the so-called *Cramer-Lundberg model* from the theory of ruin processes (Asmussen and Albrecher 2010). We note that the use of the exponential distribution will be important for the special cases we study; but our results on general distributions extend to an essentially arbitrary gap distribution G_i.

In summary, the agent's reserve at time t, denoted $R_i(t)$, is thus given by the equation

$$R_i(t) = u_i + c_i t - \sum_{j:T_{ij} \leq t} S_{ij}$$

where the last term is simply the total size of all shocks that have arrived by time t. (The number of shocks arriving by any finite time t is finite in the model with probability 1.)

Ruin Probability. Our goal is to help agents keep their reserve from becoming negative; if $R_i(t) < 0$ at any time t, then we say the agent has experienced *ruin*. Let ψ_i be the probability that there exists a time t at which agent i experiences ruin; since this is a function determined entirely by the agent's income c_i, initial reserve u_i, arrival rate of shocks β_i, and shock-size distribution F_i, we can write it as $\psi_i = \psi(c_i, u_i, \beta_i, F_i)$.

The qualitative behavior of the ruin probability ψ_i depends heavily on a parameter called *drift*, which captures the expected change per unit time in the agent's reserve. Specifically, if the expected value of the shock-size distribution is μ_i, then the drift is equal to $c_i - \beta_i \mu_i$. A standard result is that if the drift is negative — so that the agent's reserve is being pulled downwards in expectation — then the ruin probability ψ_i is equal to 1: the agent will be ruined almost surely as t goes to infinity. On the other hand, if the drift is positive, there is still a chance that agent can be ruined by shocks that are large and/or rapid enough; but it can be shown that $\psi_i < 1$, so there is a positive probability that the agent will never be ruined even as time runs to infinity (Asmussen and Albrecher 2010). In the special cases we study, we focus on the case of positive drift, where the agents might be able to avoid ruin on their own, but we would like to help lower their ruin probabilities. When we move to the general case, we will allow both positive and negative drift.

Optimization. We now consider how to model the problem of providing assistance to the agents. Let us first consider the case of *income subsidies*, in which we have a budget B, and we can choose to increase the income of agent i by an amount x_i, as long as the total amount $\sum_{i=1}^{n} x_i$ that is given out is at most B. We would like to do this so as to reduce some objective function based on the ruin probabilities of the agents. The choice of objective function reflects a societal preference on which outcomes are most desirable; to be concrete, we observe that two natural objectives are a *min-sum* formulation and a *min-max* formulation.

In the min-sum formulation, each agent i has a weight w_i representing the social cost resulting from ruin to agent i. The min-sum objective seeks to minimize the weighted expected number of agents who experience ruin:

$$\min_{x_1+\cdots+x_n=B} \sum_{i=1}^{n} w_i \psi(c_i + x_i, u_i, \beta_i, F_i),$$

where we observe that $\psi(c_i + x_i, u_i, \beta_i, F_i)$ denotes the ruin probability of agent i after a subsidy of x_i has been added to their income.

In contrast, the min-max formulation seeks to ensure that the worst ruin probability experienced by any agent is as low as possible:

$$\min_{x_1+\cdots+x_n=B} \max_{i=1,\ldots,n} \psi(c_i + x_i, u_i, \beta_i, F_i).$$

These functions correspond to two well-studied societal objectives. Of course, these are not the only two reasonable objective functions. Societal implications in this choice are discussed in the appendix.

Instead of an income subsidy, we could alternatively consider a *wealth subsidy*: an amount z_i is added to agent i's initial reserve so as to reduce the ruin probability. We can again formulate min-sum and min-max versions of the problem with wealth subsidies; the only difference at the level of the notation set-up is that the ruin probability of agent i with wealth subsidy z_i is evaluated as $\psi(c_i, u_i + z_i, \beta_i, F_i)$.

3 Agents with Zero Initial Reserve

We begin with a special case of our problem, which will shed light on some of the main techniques of subsequent results. We consider this first for the fundamental case of agents who have no initial reserve. This is a natural instance of the problem to explore given that the empirical and policy work in these domains generally focuses on instances where individuals have almost no existing buffer against ruin. Assuming an initial reserve equal to 0 is an abstraction of this challenging case.

As before, each agent i is characterized by an income c_i, a reserve u_i, which is equal to 0 in the present case, and shocks that arrive according to a Poisson process of rate β_i, and with sizes drawn from a distribution F_i. We are interested in the probability that agent i will experience eventual ruin; this is given by $\psi_i = \psi(c_i, 0, \beta_i, F_i)$.

As is standard in the theory of ruin processes, we will make the following mild assumption about the shock-size distribution F_i throughout the paper — that if we let Z_t denote the random variable equal to the total size of all shocks occurring between times 0 and t, then the quantity Z_t/t (the average amount of shock per unit time) converges to a constant limit with probability 1. This condition is satisfied whenever the shock-size/shock-arrival distributions have finite mean and variance, and therefore essentially all distributions we might wish to consider.

A fundamental result in the theory of ruin processes is that when the initial reserve is 0, the ruin probability ψ depends on the shock-size distribution F_i only through its mean value μ_i: if F_i and F_i' are shock-size distributions with the same means, then $\psi(c_i, 0, \beta_i, F_i) = \psi(c_i, 0, \beta_i, F_i')$. Thus, in a mild extension of our notation, we will write $\psi(c_i, 0, \beta_i, \mu_i)$ to stand in for $\psi(c_i, 0, \beta_i, F_i)$, when μ_i is the mean of F_i. Moreover, the ruin probability has a particular clean functional form

$$\psi(c_i, 0, \beta_i, \mu_i) = \frac{\beta_i \mu_i}{c_i}. \tag{1}$$

Here, we will focus on income subsidies. To analyze this process in terms of the min-max and min-sum objectives, it is useful to formulate the underlying optimization problem more abstractly. This abstraction will be useful for other special cases we consider as well as the generalized version of our problem. We do this next, before returning to the application for agents with zero initial reserve.

An Abstract Formulation

There is an abstract optimization problem that will provide a useful unifying description for the current problem and several of the subsequent ones we consider. The problem and its solution are related to "water-filling algorithms" from the theory of convex minimization (Alaei et al. 2012; Boyd and Vandenberghe 2004); we describe it here because the form of the solution provides an important structural insight for our domain — that the optimal allocation of subsidies in each case is based on a *priority ordering* of the agents. We first describe this abstract problem and its solution, and then we show how it applies to agents with zero initial reserve. Proofs can be found in the appendix. Our problem is as follows:

($*$) We have functions $f_1, \ldots, f_n$. Each f_i is a continuous function of a single real-variable that is positive and strictly decreasing: if $x < z$ then $f_i(x) > f_i(z)$. We would like to find non-negative real numbers $x_1, \ldots, x_n$ so that $\sum_{i=1}^{n} x_i = B$ and $\max_i f_i(x_i)$ is minimized.

Intuitively, if we think of f_i as the ruin probability of agent i when given income subsidy x_i, we see that the min-max objective is a direct special case of Problem ($*$). But, as we will see later in the section, this formulation will allow us to solve the min-sum objective as well. It is useful to first discuss the special case where all $f_i(0)$ are the same.

Lemma 1. *If $f_i(0) = f_j(0)$ for all i and j, then there is a unique vector $x^* = (x_1^*, \ldots, x_n^*)$ with the property that $\sum_i x_i^* = B$ and $f_i(x_i^*) = f_j(x_j^*)$ for all i and j. Moreover, x^* uniquely optimizes Problem $(*)$ in this case.*

Building on this special case, we would like to study the behavior of the following *priority algorithm* for solving Problem $(*)$. First, we index the functions so that $f_1(0) \geq f_2(0) \geq \cdots \geq f_n(0)$; that is, if $i \leq j$, then $f_i(0) \geq f_j(0)$. The algorithm is then easy to describe informally: we increase x_1 continuously until $f_1(x_1)$ matches $f_2(0)$; we then continuously increase both x_1 and x_2 simultaneously, keeping the values of $f_1(x_1)$ and $f_2(x_2)$ equal to each other, until they both match $f_3(0)$; we then continuously increase x_1, x_2, x_3 simultaneously, keeping their values equal to each other, until they all match $f_4(0)$; and, we proceed in this way until we reach the budget B. A more formal description of the algorithm can be found in the appendix. We think of this as a *priority algorithm* because it arrives at a solution by increasing the values of x_i in a natural priority ordering: first just x_1, then both x_1 and x_2 simultaneously, and so forth.

We now establish the basic properties of the solution returned by this priority algorithm.

Theorem 2. *Let $x^* = (x_1^*, \ldots, x_n^*)$ be the solution returned by the priority algorithm. Then x^* is the unique vector that minimizes the objective function $\max_i f_i(x_i)$, and it is also the unique vector satisfying the following property:*

(†) (i) If x_i and x_j are both positive, then $f_i(x_i) = f_j(x_j)$; and (ii) if $x_i > 0$ but $x_j = 0$, then $f_i(x_i) \geq f_j(x_j)$.

The Case of Zero Initial Reserve

We now return to our motivating question — how to optimally allocate income subsidies to agents with zero initial reserve. As before, the ruin probability of agent i, with income c_i and shocks of arrival rate β_i and mean size μ_i, is given by $\psi(c_i, 0, \beta_i, \mu_i) = \frac{\beta_i \mu_i}{c_i}$.

For the min-max objective, we can directly apply the priority algorithm developed for the formulation in the preceding subsection. We define $f_i(x_i) = \psi(c_i + x_i, 0, \beta_i, \mu_i) = \frac{\beta_i \mu_i}{c_i + x_i}$, and we find a subsidy to minimize $\max_i f_i(x_i)$. The priority ordering used by the algorithm in this case is the ruin probability itself, $f_i(0) = \psi(c_i, 0, \beta_i, \mu_i)$.

For the min-sum objective, we will also be able to use the abstract formulation as follows. Let the objective function in the min-sum case by denoted by a function ϕ where

$$\phi(x_1, \ldots, x_n) = \sum_{i=1}^{n} w_i \psi(c_i + x_i, 0, \beta_i, \mu_i).$$

Using Equation 1, we can restate this to be

$$\phi(x_1, \ldots, x_n) = \sum_i w_i \frac{\beta_i \mu_i}{c_i + x_i}.$$

We now take the partial derivative of ϕ with respect to x_i:

$$\frac{\partial \phi}{\partial x_i} = -\frac{w_i \beta_i \mu_i}{(c_i + x_i)^2}. \quad (2)$$

We see that ϕ is strictly convex by taking the second partial derivative with respect to x_i

$$\frac{\partial^2 \phi}{\partial^2 x_i} = \frac{2 w_i \beta_i \mu_i}{(c_i + x_i)^3}$$

which is strictly positive for all $x_i \geq 0$.

Because of the strict convexity, ϕ has a unique local (and hence also global) minimum over the set defined by $x_i \geq 0$ and $\sum_i x_i = B$. We can characterize it using the priority algorithm from our abstract formulation as follows. We define $f_i(x_i) = -\frac{\partial \phi}{\partial x_i}$, and we find the $x^* = (x_1^*, \ldots, x_n^*)$ that minimizes $\max_i f_i(x_i)$. By Theorem 2, the resulting point x^* has the property that $\frac{\partial \phi}{\partial x_i} = \frac{\partial \phi}{\partial x_j}$ whenever both x_i and x_j are positive, and $\frac{\partial \phi}{\partial x_i} \leq \frac{\partial \phi}{\partial x_j}$ when $x_i > 0$ and $x_j = 0$. (Recall that the partial derivatives are all negative, so f_i in our application of the abstract formulation is the negative of the corresponding partial derivative.) This implies that x^* is a local minimum for ϕ and thus, by strict convexity, it is the unique global minimum.

By using the priority algorithm and Theorem 2, we see that the choice of agents to assist with subsidies in the min-sum case proceeds via a priority rule, but one that is based neither on income nor on ruin probability. Rather, the priority given to agents is based on $\frac{\partial \phi}{\partial x_i}$ evaluated at 0

$$f_i(0) = \frac{w_i \beta_i \mu_i}{c_i^2}.$$

Interestingly, since the ruin probability is $\frac{\beta_i \mu_i}{c_i}$, it follows that the priority is, in fact, the product of three terms: the agent's weight, their ruin probability, and the reciprocal of their income.

Contrasting Prioritizations and Efficiency Loss

For both the min-max and min-sum objectives, we have seen a way to optimally allocate income subsidies in settings where agents have no initial wealth. Note that these results hold for any general distribution from which the shocks may be drawn and only require the mean size of the distribution. Another key takeaway is that these algorithms inherently propose a *priority ordering* of the agents by need for the given objective. We can therefore ask: how different can these orderings be from one another? What about the ordering that simply uses the agents' income?

The three priority orderings under consideration are by: (I) income, (II) ruin probability, and (III) the ordering used for the min-sum objective. We investigate how different these priority orderings can be from one another.

Lemma 3. *The priority orderings given by (I) income and (III) priority orderings given by Equation 2 (with unit weights) can be the reverse of one another. Likewise, the priority orderings given by (II) ruin probability can be the reverse of (III) priority orderings given by our solution.*

The proof of this lemma can be found in the appendix. On the other hand, we find that there are some dependencies among the priority orderings, including the following immediate consequence.

Corollary 4. If the priority ordering by (I) income and (II) probability of ruin are the same, then the priority ordering (III) given by Equation 2 for the min-sum objective (with unit weights) will also be the same as this ordering.

These results show that the priority ordering can be different depending on whether we use income, ruin probability, or our optimal solution. It is natural to ask what the potential cost of using a more naive prioritization scheme would be, and we find that it can be high. Namely, we find that there can be a gap of $\Omega(\sqrt{n})$ comparing our solution with both income and ruin probability. See appendix.

The Exponential Case. The case of non-zero wealth becomes much more complex in general; for example, we do not have closed-form expressions for the ruin probability with general distributions as we do for the case of zero wealth. To get a sense for the properties of non-zero wealth in a setting that has complex behavior but still exhibits closed-form solutions, we consider the case when shocks are drawn from an exponential distribution F_i. We pose questions comparing priority orderings for solutions for the income subsidy and wealth subsidy problems for both the min-sum and min-max objectives. The appendix contains a full discussion of results showing a range of counter-intuitive observations about priority orderings and monotonicity of the quantities of interest in the agents' parameters.

4 Non-zero Wealth: General Distributions

The previous sections characterize optimal subsidies when all agents have positive drift, together with either zero initial wealth or shock distributions satisfying a specific functional form. We now consider the case of general shock distributions, arbitrary initial wealth, and arbitrary drift. This section contains three results:

- Lemma 5: A polynomial-time algorithm for the min-max objective when agents have arbitrary shock distributions, initial wealth, and drift.
- Theorem 6 (main result): A Fully Polynomial-Time Approximation Scheme (FPTAS) for the min-sum objective when agents have arbitrary shock distributions, initial wealth, and drift.
- Proposition 7: A proof that the min-sum objective is (weakly) NP-hard in general, implying that we should not expect better than an FPTAS without further assumptions on the problem.

Min-Max via Binary Search

We first show how to allocate income subsidies so as to optimize the min-max objective. Intuitively, our algorithm looks like a binary search for the optimal min-max value X. For each potential guess p of X, we just need to see, for each agent, how much income subsidy is required to achieve a ruin probability p. If the sum of these subsidies exceeds B, then it is infeasible to have min-max value p. Otherwise, it is feasible. So, we can repeatedly guess p in binary search and converge quickly to the optimum. Lemma 5 essentially formalizes this intuition while being careful about the cost of certain operations. For example, throughout this section, we will assume for simplicity of exposition that $\psi^{-1}_{u,\beta,F}(\cdot)$ can be computed in $O(1)$ operations, where $\psi_{u,\beta,F}(c) = \psi(c, u, \beta, F)$ (that is, the minimum x such that $\psi(x, u, \beta, F) \leq p$ can be computed in $O(1)$ operations for all p). We briefly discuss at the end of this section ways in which this assumption can be relaxed.

Below, we seek to allocate income subsidies $x_1, \dots, x_n$ with $\sum_i x_i = B$ such that agent i receives subsidy x_i in a way that minimizes $\max_i\{\psi(c_i + x_i, u_i, \beta_i, F_i)\}$.

Lemma 5. *Let X denote the optimum value for the min-max objective for any given instance. Then, for any $\delta > 0$, a solution with min-max value $X + \delta$ can be found in time* $\text{poly}(n, \log(1/\delta))$.

Proof. The algorithm is based on binary search. We first need a subroutine to check, for a given value p, whether it is feasible to subsidize all agents to probability of ruin at most p. To this end, simply compute $\psi^{-1}_{u_i,\beta_i,F_i}(p)$ for all i, and update $x_i := \max\{0, \psi^{-1}_{u_i,\beta_i,F_i}(p) - c_i\}$. If $\sum_i x_i \leq B$, then these choices of x_i explicitly witness that it is feasible to have min-max objective $\leq p$. If not, then they explicitly prove that with a budget constraint of B, some agent must have ruin probability exceeding p.

It therefore suffices to do binary search on the interval $[0, 1]$. Each iteration of binary search takes $O(n)$ operations, and therefore doing $\log(1/\delta)$ iterations of binary search takes $O(n \log(1/\delta))$ operations. After $\log(1/\delta)$ iterations, we will have a window $[X, X+\delta]$ where we explicitly found a choice of subsidies guaranteeing min-max objective $X+\delta$, and also proved that better than X is not possible. Thus, we have an additive δ approximation in the desired time. ☐

Min-Sum via Knapsack

We now show how to make use of our min-max approximation as a subroutine to provide an FPTAS for the min-sum objective. Our approach barely uses any structure of the $\psi_{u_i,\beta_i,F_i}(\cdot)$ function, and is essentially an FPTAS to minimize $\sum_i f_i(x_i)$ subject to $\sum_i x_i \leq B$ for any no-increasing functions $f_i(\cdot)$. Our algorithm is also very similar to the FPTAS for Knapsack via dynamic programming.

Theorem 6. *Let X denote the optimum value for the min-sum objective for any given instance. Then, for any $\varepsilon, \delta > 0$, a solution with min-sum value $(1+\varepsilon)X + \delta$ can be found in time* $\text{poly}(n, 1/\varepsilon, \log(1/\delta))$.

Proof. The algorithm proceeds in two phases. First, we need to figure out the scale of the optimum solution. Then, we achieve an additive approximation at the appropriate scale by adapting the FPTAS for the weighted knapsack problem (Vazirani 2013).

For phase one, we simply solve the min-max objective using Lemma 5 up to accuracy $\delta/(2n)$, in time $\text{poly}(n, \log(n/\delta))$. Let C denote the value of the solution output by the min-max subroutine. If $C \leq \delta/n$, then certainly the min-sum objective for this same solution is at most δ, which satisfies the desired guarantee (as $X \geq 0$). Otherwise, we know from Lemma 5 that it is certainly not possible to get the ruin probabilities to sum to less than $C - \delta/(2n)$ (since, otherwise, the max would be at most $C - \delta/(2n)$ as well, which would contradict Lemma 5). As $C \geq \delta/n$, this means that the optimum is at least $C/2$. Therefore, if we set $\eta := \varepsilon C/(2n)$ and get an additive $n\eta$ approximation, this will be a multiplicative $(1+\varepsilon)$ approximation.

Observe also that our min-max subroutine explicitly outputs a solution with min-sum value at most Cn (because the maximum ruin probability is at C). We thus know that the optimum lies somewhere between $C/2$ and Cn. These bounds are sufficient to use a dynamic program similar to that for Knapsack.

To create our Dynamic Program, first define:

$$f_i(x_i) := \lceil \psi(c_i + x_i, u_i, \beta_i, F_i)/\eta \rceil \cdot \eta.$$

That is, $f_i(x_i)$ is $\psi(c_i + x_i, u + i, \beta_i, F_i)$ rounded up to the nearest integer multiple of η. Note that this only makes any potential solution worse by at most an additive factor of η. After this rounding, we claim we can find the optimal solution with dynamic programming. Briefly observe that $f_i^{-1}(\cdot)$ can be computed with one black-box call to $\psi^{-1}_{u_i,\beta_i,F_i}(\cdot)$ (which are still assuming can be computed in $O(1)$ operations for ease of exposition).

Let $G(j,k)$ denote the minimum possible budget that suffices to guarantee $\sum_{i \leq j} f_i(x_i) = k \cdot \eta$. We claim this can be found using the recurrence below.

$$G(j,k) = \min_{\ell}\{G(j-1, k-\ell) + f_i^{-1}(\ell\eta)\},$$

where ℓ ranges in $\{0, 1, \dots, k\}$.

Observe first that the range of ℓ suffices, as all values are integer multiples of η. Observe also that the optimal solution for $G(j,k)$ must obtain some value $\ell\eta$ for $f_j(x_j)$, which means it must get $(k-\ell)\eta$ from the first $j-1$, and that the optimization for the first $j-1$ is exactly covered by $G(j-1, k-\ell)$. Finally, observe that each step requires taking a minimum over at most k terms, and that the entire DP table has $n * K$ terms, if we let k range in $\{0, \dots, K-1\}$. The entire table can thus be filled in time $\text{poly}(n, K)$. We will set $K := n^3/\varepsilon$.

We would like to find the minimum k such that $G(n,k) \leq B$. Observe that as $K := n^3/\varepsilon$, it is the case that $G(n,K) \leq B$, as this corresponds to a solution of value $n^3 \cdot \eta = Cn^2$, which is guaranteed to exist by our work in phase one. Observe that this dynamic program finds an optimal allocation, up to an additive $n\eta = \varepsilon C/2$. So, as the optimum is at least $C/2$, this is at most $(1+\varepsilon)$OPT, as desired. □

NP-Hardness and Further Computational Considerations

We prove an NP-hardness result showing that one should not expect an exact solution in poly-time without some assumptions, and also discuss relaxations of the assumption that $\psi^{-1}_{u,\beta,F}(\cdot)$ can be computed in $O(1)$ operations.

First, recall that the approach used in Theorem 6 is quite general: it essentially provides an FPTAS for minimizing $\sum_i f_i(x_i)$ subject to $\sum_i x_i \leq B$ for *any* non-increasing functions such that $f_i^{-1}(\cdot)$ can be computed efficiently. We first show that this problem is NP-hard in general, even for fairly simple functional forms of $f_i(\cdot)$.

Proposition 7. Let MIN-SUM take as input explicit descriptions of n functions $f_i(\cdot)$, and output $\min_{\vec{x}, x_i \geq 0 \forall i, \sum_i x_i \leq B}\{\sum_i f_i(x_i)\}$. Then, MIN-SUM is (weakly) NP-hard, even when each $f_i(\cdot)$ takes the form of $\min\{1, g_i(\cdot)\}$, where $g_i(\cdot)$ is convex.[2]

It is certainly possible that for certain functional forms of F_i (e.g., exponentially distributed), that the min-sum objective can be optimized exactly in polynomial time. But the above hardness shows that and FPTAS is likely the best one should hope for without further assumptions. Finally, we briefly discuss further computational considerations with regards to computing $\psi^{-1}_{u_i,\beta_i,F_i}(\cdot)$. Proofs can be found in the corresponding section in the appendix.

Lemma 8. *Let $f(\cdot)$ be non-increasing, and have a domain of $[0, B]$. Then if $f(\cdot)$ can be computed in $O(1)$ operations, for all x, a y satisfying $y \leq f^{-1}(x) \leq y+\delta$ can be computed in $O(\ln(B/\delta))$ operations.*

Proof. Let the input query be x. If $f(B) > x$, then the inverse is undefined. If $f(0) < x$, then the inverse is also undefined. Otherwise, an inverse exists in $[0, B]$. From here, we proceed with binary search. After $\log_2(B/\delta)$ steps, we will have a window of the form $[y, y+\delta]$ where $f^{-1}(x)$ certainly lies in this window. So output this y. □

This lemma alone does not transparently affect the approximation guarantees — the issue is that perhaps a little error in $y \approx f^{-1}(x)$ may cause $f(y) \gg x$. With one extra assumption on f, however (a Lipschitz condition), this is useful. Below, we say that a function $f(\cdot)$ is (L,δ)-Lipschitz if for all x, y with $|x-y| \leq \delta/L$, $|f(x) - f(y)| \leq \delta$.

Corollary 9. Let $\psi_{u_i,\beta_i,F_i}(\cdot)$ be (L,δ)-Lipschitz. Then if $\psi_{u_i,\beta_i,F_i}(\cdot)$ can be computed in $O(1)$ operations, an additive $O(\delta)$-approximation to the min-max objective can be found in time $\text{poly}(n, \log(LB/\delta))$, and a multiplicative $(1+\varepsilon)$, additive δ approximation to the min-sum objective can be found in time $\text{poly}(n, 1/\varepsilon, \log(LB/\delta))$.

Proof Sketch. Using Lemma 8, a y satisfying $y \leq \psi^{-1}_{u_i,\beta_i,F_i}(x) \leq y + \delta/L$ can be found in $O(\ln(LB/\delta))$ operations. Therefore, this value of y also satisfies $f(y) \leq x + \delta$, as ψ_{u_i,β_i,F_i} is (L,δ)-Lipschitz. Using this instead of a true inverse oracle accumulates an additional additive error of δ in both proofs, which can be accommodated into the existing guarantees. □

Intuitively, the need for *some* Lipschitz condition on $\psi_{u_i,\beta_i,F_i}(\cdot)$ (if we only wish to have black-box access to $\psi_{u_i,\beta_i,F_i}(\cdot)$ and not $\psi^{-1}_{u_i,\beta_i,F_i}(\cdot)$) is because if ψ_{u_i,β_i,F_i} jumps instantaneously from (say) 1 to 0, it is impossible to detect exactly where. Thus, in order to ensure we are competitive with the optimum, we would also need to violate the budget constraint by a little bit. Even without any Lipschitz condition, this approach suffices:

[2] Observe that without the minimum with 1, this is just convex minimization, which can be solved in polynomial time.

Lemma 10. *If all $\psi_{u_i,\beta_i,F_i}(\cdot)$ can be computed in $O(1)$ operations. Then if X denotes the optimal solution to the min-max objective with budget B, a solution with min-max value $X + \delta$ using budget $B + \delta$ can be found in time* $\mathrm{poly}(n, \log(B/\delta))$. *If Y denotes the optimal solution to the min-sum objective with budget B, a solution with min-sum value $(1 + \varepsilon) \cdot Y + \delta$ using budget $B + \delta$ can be found in time* $\mathrm{poly}(n, 1/\varepsilon, \log(B/\delta))$.

5 Related Work

The multi-faceted nature of economic welfare creates challenges for measuring and analyzing it in a tractable manner. Some references surveying related work include (Atkinson 2015; Grusky 2018). A key point in this area has been to introduce measurements that account for the complexity while preserving enough simplicity to inform policy and interventions aimed at mitigating poverty and economic inequality. In this paper, we study one dimension in this space — the phenomenon of income shocks — whose impact on economic welfare has been extensively documented through a long line of empirical work. Income shocks have been shown to have complex interactions with many factors and have long-term and severe consequences for families including eviction, job loss, and poor health outcomes (Atake 2018; O'Flaherty 2009; Shapiro and others 2004). Yet, they do not play a correspondingly central role in many social assistance programs. For a discussion on the role of income shocks on the well-being of families in New York City, see reports by the Poverty Tracker (Wimer et al. 2014).

There is a line of modeling work in the study of household consumption and public economics that considers agent-level behavior in response to subsidies and more broadly understanding consumption dynamics (Aiyagari 1994; Carroll and Samwick 1997; Zeldes 1989). We take a different approach here, focusing on the algorithmic issues inherent in the stochastic optimization problem at the heart of subsidy allocation with shocks. Our model is based on the framework of *ruin probabilities* from the literature on stochastic risk modeling (Asmussen and Albrecher 2010). This work also has close connections with stochastic control theory. There has been work related to investment in risky assets by insurance companies seeking to minimize the probability of ruin (Azcue and Muler 2009; Hipp and Plum 2000; Schmidli and others 2002). To our knowledge there has not been work looking at our problem of allocating a fixed budget over multiple agents in order to change premium flow or initial reserve.

Our paper can be viewed as belonging to an emerging style of work that uses computational and optimization-based methods to inform assistance programs aimed at improving access to opportunity for vulnerable populations. One recent instance studies allocating interventions for homelessness services using a mixture of counter-factual reasoning and mechanism design (Kube, Das, and Fowler 2018); another studies optimal allocation of financial aid in US colleges based on students' parental income (Findeisen and Sachs 2016). Relative to these recent papers, our work is the first that we are aware of to take a computational approach to assistance in a setting where the underlying optimization problem exhibits the type of rich stochastic dynamics over time that characterizes our domain.

6 Conclusion and Future Directions

We propose a model of economic welfare that incorporates agents' income and initial reserve as well as income shocks, and we analyze the model using results from the theory of ruin processes. We consider a problem faced by a planner who would either like to maximize the welfare of the most vulnerable agent or minimize the number of agents that experience ruin. Our analysis reveals several insights into the role of income shocks on economic welfare. For instance, we find that agents may appear to be less vulnerable when considering welfare measures that simply use income than measures we introduce in this paper. And, in fact, we also find that even measures proposed in this paper — i.e., ruin probability or those given by our solutions for income or wealth subsidies — can be drastically different from one another. Therefore, care must be taken in picking the desired objective as well as the type of subsidy to be given out.

The different forms of subsidies that we consider in this paper closely resemble different assistance programs or proposed subsidies. Income subsidies are reflected in programs such as the Supplemental Nutrition Assistance Program or housing vouchers which reduce families' monthly expenses, thereby leaving more reserves for families month-to-month. (In our model, this roughly corresponds to adding an income subsidy x_i to c_i.) Wealth subsidies may resemble proposed policies such as "baby bonds" or inheritance taxation to alleviate racial wealth inequalities (Hamilton and Darity Jr 2010; Shapiro and others 2004)

There are various assumptions about our model that warrant further discussion and also leave open directions for further work. We present a brief overview below and further discussions can be found in the appendix. We treat income as a steady stream. We incorporate disruptions to one's income as a form of shock. Cases in which an individual loses their job entirely, or experiences some other shock that essentially amounts to a long-term state change, may benefit from a different model that more directly incorporates such discrete transitions.

We also assume that income and initial reserves are known, well-defined values. In practice, it may be difficult to quantify these, and especially reserves. For instance, reserves such as savings may be simple to represent, whereas other forms of reserves such as social capital or having "outside options" such as ability to rely on family for certain types of shocks may be harder to quantify. Note that we address some of this via the weighted version of our problem, with potentially different weights assigned to each agent, but there are many options in how best to determine these weights and the weights themselves may also be dynamic. Similarly, we assume that the agents' incomes are fixed and they do not make choices in consumption. In many settings, agents may adjust their consumption or make investments in response to such interventions. For further discussions on the societal implications of these above assumptions and other related directions, we direct the reader to the appendix.

References

Aiyagari, S. R. 1994. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics* 109(3):659–684.

Akerlof, G. A. 1978. The economics of “tagging” as applied to the optimal income tax, welfare programs, and manpower planning. *The American economic review* 68(1):8–19.

Alaei, S.; Fu, H.; Haghpanah, N.; Hartline, J.; and Malekian, A. 2012. Bayesian optimal auctions via multi-to single-agent reduction. *arXiv preprint arXiv:1203.5099*.

Alatas, V.; Purnamasari, R.; Wai-Poi, M.; Banerjee, A.; Olken, B. A.; and Hanna, R. 2016. Self-targeting: Evidence from a field experiment in indonesia. *Journal of Political Economy* 124(2):371–427.

Asmussen, S., and Albrecher, H. 2010. *Ruin probabilities*. World Scientific Publishing Co Pte Ltd.

Atake, E.-H. 2018. Health shocks in sub-saharan africa: are the poor and uninsured households more vulnerable? *Health economics review* 8(1):26.

Atkinson, A. B. 2003. Multidimensional deprivation: contrasting social welfare and counting approaches. *The Journal of Economic Inequality* 1(1):51–65.

Atkinson, A. B. 2015. *Inequality*. Harvard University Press.

Azcue, P., and Muler, N. 2009. Optimal investment strategy to minimize the ruin probability of an insurance company under borrowing constraints. *Insurance: Mathematics and Economics* 44(1):26–34.

Boyd, S., and Vandenberghe, L. 2004. *Convex Optimization*. Cambridge University Press.

Calabresi, G., and Bobbitt, P. 1984. Tragic choices.

Carroll, C. D., and Samwick, A. A. 1997. The nature of precautionary wealth. *Journal of monetary Economics* 40(1):41–71.

Desmond, M. 2016. *Evicted: Poverty and profit in the American city*. Broadway Books.

Elster, J. 1992. *Local justice: How institutions allocate scarce goods and necessary burdens*. Russell Sage Foundation.

Eubanks, V. 2018a. Automating bias. *Scientific American* 319(5):68–71.

Eubanks, V. 2018b. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin’s Press.

Findeisen, S., and Sachs, D. 2016. Optimal need-based financial aid. Technical report, Mimeo, University of Mannheim/European University Institute.

Gilman, M. E. 2011. The class differential in privacy law. *Brook. L. Rev.* 77:1389.

Grusky, D. B. 2018. *Social stratification: Class, race, and gender in sociological perspective*. Routledge.

Hamilton, D., and Darity Jr, W. 2010. Can ‘baby bonds’ eliminate the racial wealth gap in putative post-racial america? *The Review of Black Political Economy* 37(3-4):207–216.

Hipp, C., and Plum, M. 2000. Optimal investment for insurers. *Insurance: Mathematics and Economics* 27(2):215–228.

Kube, A. R.; Das, S.; and Fowler, P. J. 2018. Allocating interventions based on counterfactual predictions: A case study on homelessness services.

Marwick, A.; Fontaine, C.; and boyd, d. 2017. “nobody sees it, nobody gets mad”: Social media, privacy, and personal responsibility among low-ses youth. *Social Media+ Society* 3(2):2056305117710455.

Mirrlees, J. A. 1971. An exploration in the theory of optimum income taxation. *The review of economic studies* 38(2):175–208.

Mkandawire, T. 2005. *Targeting and universalism in poverty reduction*. United Nations Research Institute for Social Development Geneva.

Nichols, A. L., and Zeckhauser, R. J. 1982. Targeting transfers through restrictions on recipients. *The American Economic Review* 72(2):372–377.

O’Flaherty, B. 2009. What shocks precipitate homelessness. *Unpublished manuscript. New York: Columbia University*.

Peppet, S. R. 2011. Unraveling privacy: The personal prospectus and the threat of a full-disclosure future. *Nw. UL Rev.* 105:1153.

Persad, G.; Wertheimer, A.; and Emanuel, E. J. 2009. Principles for allocation of scarce medical interventions. *The Lancet* 373(9661):423–431.

Roth, A. E. 2007. Repugnance as a constraint on markets. *Journal of Economic perspectives* 21(3):37–58.

Roth, A. E. 2008. What have we learned from market design? *Innovations: Technology, Governance, Globalization* 3(1):119–147.

Schmidli, H., et al. 2002. On minimizing the ruin probability by investment and reinsurance. *The Annals of Applied Probability* 12(3):890–907.

Shapiro, T. M., et al. 2004. *The hidden cost of being African American: How wealth perpetuates inequality*. Oxford University Press, USA.

Skocpol, T. 1991. Targeting within universalism: Politically viable policies to combat poverty in the united states. *The urban underclass* 411(411):437–59.

Vazirani, V. V. 2013. *Approximation algorithms*. Springer Science & Business Media.

Wimer, C.; Bannerman, C.; Garfinkel, I.; and Neckerman, K. 2014. Poverty tracker – monitoring poverty and well-being in NYC: Report 2. Technical report, Columbia Population Research Center and Robin Hood.

Zeldes, S. P. 1989. Consumption and liquidity constraints: an empirical investigation. *Journal of political economy* 97(2):305–346.

A An Abstract Optimization Problem

We first formally describe the priority algorithm from the main text.

- For each $i \leq j$, we find the value δ_{ij} such that $f_i(\delta_{ij}) = f_j(0)$: that is, the input we can provide to f_i if we want to reduce its value to $f_j(0)$. Since each function is strictly decreasing, there is a unique such δ_{ij}.
- We now find the maximum j such that $\sum_{i<j} \delta_{ij} \leq B$. Let m be this value of j; note that $m \geq 1$. (For notational simplicity, we define $\delta_{mm} = 0$.)
- We are now at a value of m such that if we tried to increase each of $x_1, \ldots, x_m$ so that $f_i(x_i) = f_{m+1}(0)$, we would exceed our budget B.
- If $\sum_{i<m} \delta_{im} = B$, then we declare this to be our solution: we set $x_i = \delta_{im}$ for $i < m$, and $x_i = 0$ for $i \geq m$.
- Otherwise, if $\sum_{i<m} \delta_{im} = C < B$, then we apply Lemma 1 as follows. We define a function $g_i(z_i) = f_i(\delta_{im} + z_i)$ for $i \leq m$, and then we apply Lemma 1 to the functions $g_1, \ldots, g_m$ with budget $B - C$. This produces a unique vector $z^* = (z_1^*, \ldots, z_m^*)$. As our solution, we report $x_i = \delta_{im} + z_i^*$ for $i \leq m$, and $x_i = 0$ for $i > m$.

We can now present proofs for Lemma 1 and Theorem 2 from the section on an abstract formulation of our optimization problem here:

Lemma 1. If $f_i(0) = f_j(0)$ for all i and j, then there is a unique vector $x^* = (x_1^*, \ldots, x_n^*)$ with the property that $\sum_i x_i^* = B$ and $f_i(x_i^*) = f_j(x_j^*)$ for all i and j. Moreover, x^* uniquely optimizes Problem $(*)$ in this case.

Proof. First, since $\max_i f_i(x_i)$ is a continuous function on the compact set defined by $\sum_i x_i = B$ and $x_i \geq 0$ for all i, it achieves its minimum at some vector $x^* = (x_1^*, \ldots, x_n^*)$.

Let q be the common value of all $f_i(0)$. We observe that the optimal value of $\max_i f_i(x_i)$ must be strictly less than q, since setting $x_i = B/n$ satisfies $\max_i f_i(x_i) < q$ by the strict decreasing property of all f_i. Thus, we must have $x_i^* > 0$ for all i, since otherwise would have $\max_i f_i(x_i^*) = q$, contradicting the optimality of x^*.

We claim that x^* satisfies the desired condition $f_i(x_i^*) = f_j(x_j^*)$ for all i and j. For, if not, let S be the set of indices i for which $f_i(x_i^*)$ is maximal. Since we do not have $f_i(x_i^*) = f_j(x_j^*)$ for all i and j, there is some index h that does not belong to S. But now, by the continuity and strict decreasing property of the functions f_i, we can choose a sufficiently small $\varepsilon > 0$, increase x_i^* by ε for each $i \in S$, and (since $x_h^* > 0$) decrease x_h^* by $\varepsilon|S|$, so as to reduce the value of $\max_i f_i(x_i)$. This contradiction shows that we must have $f_i(x_i^*) = f_j(x_j^*)$ for all i and j.

To show uniqueness, suppose there were two distinct vectors x^* and z^* that both satisfied the equality $f_i(x_i^*) = f_j(x_j^*)$ and $f_i(z_i^*) = f_j(z_j^*)$ for all i, j. Then, for some i, we must have $x_i^* \neq z_i^*$; suppose (by symmetry) that $x_i^* > z_i^*$. Since f_i is strictly decreasing, it follows that $f_i(x_i^*) < f_i(z_i^*)$. Now, since all f_j are strictly decreasing, we would then have $f_j(x_j^*) < f_j(z_j^*)$ for all j, from which it follows that $x_j^* > z_j^*$ for all j. But, this contradicts the assumption that both $\sum_i x_i^*$ and $\sum_i z_i^*$ are equal to B. □

Theorem 2. Let $x^* = (x_1^*, \ldots, x_n^*)$ be the solution returned by the priority algorithm. Then x^* is the unique vector that minimizes the objective function $\max_i f_i(x_i)$, and it is also the unique vector satisfying the following property:

(†) (i) If x_i and x_j are both positive, then $f_i(x_i) = f_j(x_j)$; and (ii) if $x_i > 0$ but $x_j = 0$, then $f_i(x_i) \geq f_j(x_j)$.

Proof. Let $f^* = \max_i f_i(x_i^*)$. If $v = (v_1, \ldots, v_n)$ is any vector whose coordinates sum to B but is not equal to x^*, then we have $v_i < x_i^*$ for some index i. Let m be the index computed by the priority algorithm. The vector x^* only has positive coordinates for indices $h \leq m$, and $f_h(x_h^*) = f^*$ for all such indices. Since f_i is strictly decreasing, we thus have $f_i(v_i) > f_i(x_i^*) = f^*$, and hence v cannot be optimal. Thus, x^* is the unique vector that minimizes the objective function.

It must also hold that $v_j > x_j^*$ for some index j. Then, $f_j(v_j) < f^*$. Since x^* is optimal, there must also be an index i for which $f_i(v_i) \geq f^*$. Either $v_i > 0$, which violates part (i) of (†), or $v_i = 0$, which violates part (ii) of (†). Thus, x^* is the unique vector that satisfies (†). □

B The Case of Zero Initial Reserve

We present missing proofs and discussions from the main text below.

We first present proof of Lemma 3, which shows contrasting prioritization between the priority orderings by (I) income, (II) ruin probability, and (III) the priority ordering used for the min-sum objective. Note, since the weights w_i used in the min-sum objective are a potential source of difference in these orderings for reasons independent of any of the other parameters, it is most interesting for the purposes of this question to consider the case in which these weights do not play a role; thus, we assume for this discussion that we have *unit weights*, with $w_i = 1$ for all agents i.

Lemma 3. The priority orderings given by (I) income and (III) priority orderings given by Equation 2 (with unit weights) can be the reverse of one another. Likewise, the priority orderings given by (II) ruin probability can be the reverse of (III) priority orderings given by our solution.

Proof. For notational convenience, we first set $r_i = \beta_i \mu_i$. This value corresponds to the total amount of shock experienced per unit time. Therefore, the probability of ruin is $\psi_i(c_i, 0, \beta_i, \mu_i) = r_i / c_i$ and

$$\frac{\partial \phi}{\partial x_i} = \frac{-r_i}{(c_i + x_i)^2}.$$

We assume that we have agents with parameters:

$$((c_1, r_2), (c_2, r_2), \ldots, (c_n, r_n)).$$

Consider a set of n agents labeled $i = (1, 2, \ldots, n)$ with income $c_i = 1 + i\epsilon$ for some $\epsilon > 0$. The shock volume

r_i is chosen to equal $(0.5 + 4i\epsilon)^2$. The priority ordering by income in this setting is $(1, 2, \ldots, n)$. On the other hand, the priority ordering by $\frac{\partial \phi}{\partial x_i}$ would rank the agents in increasing order of

$$-\frac{r_i}{c_i^2} = -\frac{(0.5 + 4i\epsilon)^2}{(1 + i\epsilon)^2}.$$

This quantity decreases as i increases. Therefore, the priority ordering by Equation 2, which gives us the optimal solution for the min-sum objective, is the reverse of the priority ordering by the agents' income.

This example also gives us a case where the priority ordering by income is the reverse of the priority ordering by ruin probability. In particular, the ruin probability is given by

$$\frac{r_i}{c_i} = \frac{(0.5 + 4i\epsilon)^2}{(1 + i\epsilon)}.$$

This value increases as i increases.

Finally, to note that the priority ordering by Equation 2 and ruin probability can be the reverse of one another, we construct an example as follows: keep r_i to be the same as above, but let $c_i = 1 + 4i\epsilon$. Then, the probability of ruin will be

$$\frac{r_i}{c_i} = \frac{(0.5 + 4i\epsilon)^2}{(1 + 4i\epsilon)},$$

which is, again, increasing in i. On the other hand, the $\frac{\partial \phi}{\partial x_i}$ values are

$$-\frac{r_i}{c_i^2} = -\frac{(0.5 + 4i\epsilon)^2}{(1 + 4i\epsilon)^2}.$$

This value decreases as i increases. Therefore, the priority orderings given by the ruin probability and Equation 2 can be reverses of one another. □

We also present the example establishing the gap of $\Omega(\sqrt{n})$ comparing our solution with both income and ruin probability.

Example. Suppose we have three types of agents A, B, and C, and all agents have $\beta_i = 1, u_i = 0$. We describe each of the three types of agents below:

A: Each agent has $\mu_i = m$ and $c = \frac{11}{10}m$. There are m agents of this type.

B: This is one agent with $\mu = m^3$ and $c = m^3$.

C: Each agent has $\mu_i = 0$ and $c_i = 1$. There are m^2 agents of this type.

The ruin probabilities for agents of type A, B, and C are $10/11, 1$, and 0, respectively, and there are a total of $m^2 + m + 1$ agents. The expected number of agents experiencing ruin is $\frac{10m}{11} + 1$.

Suppose the designer has a budget of $\frac{1}{10}m^3$. We consider the three types of priority orderings:

I: A priority ordering by income would focus on type C agents. We divide up our budget evenly, and each agent would receive an income subsidy of $m/10$. Since these agents already had ruin probability 0, the expected number of agents experiencing ruin remains unchanged.

II: Next, we consider prioritizing the agents by ruin probability. In this case, we would give our entire budget to the agent of type B, reducing their ruin probability to $10/11$. In this case, the expected number of agents experiencing ruin is still $\Theta(m)$.

III: Finally, we consider our solution: we would prioritize agents of type A, dividing up our budget evenly among the agents. In this case, the ruin probability of each agent in A is reduced to $O(1/m)$, and so the expected number of agents experiencing ruin is $O(1)$.

We therefore have a gap of $\Omega(m)$ between our solution and both the priority orderings that use income or ruin probability. Since there are $n = \Theta(m^2)$ agents, the gap is proportional to the square root of the number of agents.

C Non-zero Wealth: The Exponential Case

We consider the case when shocks are drawn from an exponential distribution F_i: the probability that a shock has size exceeding x is given by $e^{-\delta x}$ for a parameter $\delta > 0$, resulting in a mean shock size of $1/\delta$. By studying this special case, we highlight counter-intuitive results showing non-monotonicity of agent's welfare in the various parameters of interest as well as rich examples where priority orderings by income, wealth, and our solutions for optimal income and wealth subsidy can vary substantially.

Specifically, suppose an agent has income c_i, initial wealth u_i, and experiences shocks at rate β_i with a size distribution that is exponential with mean $1/\delta_i$; then a result from the theory of ruin processes (Asmussen and Albrecher 2010) shows that the ruin probability is

$$\psi(c_i, u_i, \beta_i, F_i) = \frac{\beta_i}{c_i \delta_i} e^{\left(\frac{\beta_i}{c_i} - \delta_i\right) u_i}. \quad (3)$$

Since the distribution F_i is fully characterized by the parameter δ_i, we will also write this as $\psi(c_i, u_i, \beta_i, \delta_i)$.

We consider three forms of subsidies – a wealth subsidy, income subsidy, and mixed income and wealth subsidy – in this setting. Below, we show how to optimally allocate subsidies in each of the three settings.

Wealth Subsidies. A wealth subsidy of z_i reduces the ruin probability from $\psi(c_i, u_i, \beta_i, \delta_i)$ to $\psi(c_i, u_i + z_i, \beta_i, \delta_i)$. For the min-max objective, we can proceed exactly as in the previous section, applying the priority algorithm from Section 3 using the functions $f_i(z_i) = \psi(c_i, u_i + z_i, \beta_i, \delta_i)$, and hence ordering agents by their ruin probabilities.

For the min-sum objective, we use partial derivatives as in the previous section as well. Specifically, we define our objective function to be

$$\begin{aligned} \gamma(z_1, \ldots, z_n) &= \textstyle\sum_{i=1}^n w_i \psi(c_i, u_i + z_i, \beta_i, \mu_i) \\ &= \textstyle\sum_i w_i \frac{\beta_i}{c_i \delta_i} e^{\left(\frac{\beta_i}{c_i} - \delta_i\right)(u_i + z_i)}. \end{aligned}$$

We denote the term associated with z_i with γ_i. Taking the partial derivative with respect to z_i, we get

$$\frac{\partial \gamma}{\partial z_i} = \left(\frac{\beta_i}{c_i} - \delta_i\right) \frac{w_i \beta_i}{c_i \delta_i} e^{\left(\frac{\beta_i}{c_i} - \delta_i\right)(u_i + z_i)}. \quad (4)$$

Since the process has positive drift, all of the partial derivatives are negative. It is easy to verify that γ is strictly convex since the second derivative with respect to z_i is strictly positive. We therefore define $f_i(z_i) = -\frac{\partial \gamma}{\partial z_i}$ and use this as a priority ordering for the algorithm from Section 3.

Income Subsidies. The case for the income subsidy proceeds in the same way as the wealth subsidy solution above. We optimize for the min-max objective using the ruin probabilities. For the min-sum objective function, we define

$$\begin{aligned}\phi(x_1, \dots, x_n) &= \sum_{i=1}^n w_i \psi(c_i + x_i, u_i, \beta_i, \mu_i) \\ &= \sum_i w_i \frac{\beta_i}{(c_i+x_i)\delta_i} e^{\left(\frac{\beta_i}{(c_i+x_i)\delta_i} - \delta_i\right)u_i}.\end{aligned}$$

We take the partial derivative with respect to x_i to get

$$\frac{\partial \phi}{\partial x_i} = -\frac{w_i \beta_i \, \mathrm{e}^{-u_i \left(\delta_i - \frac{\beta_i}{c_i+x_i}\right)}}{\delta_i \left(c_i + x_i\right)^2} - \frac{\beta_i^2 \, u_i \, \mathrm{e}^{-u_i \left(\delta_i - \frac{\beta_i}{c_i+x_i}\right)}}{\delta_i \left(c_i + x_i\right)^3}. \tag{5}$$

We again has a convex minimization problem and we use $f_i(x_i) = -\frac{\partial \phi}{\partial x_i}$ as an ordering for the priority algorithm to find the optimal solution.

Mixed Income and Wealth Subsidies.

We can also study a mixed income and wealth subsidy, in which there is a "conversion factor" k between money allocated for wealth subsidies and income subsidies; the planner solving the optimization problem can give income subsidy x_i and wealth subsidy z_i to agent i, provided that $\sum_i z_i + k \sum_i x_i \leq B$. Note that although a wealth subsidy is implemented as a one-time allotment while an income subsidy is paid out over time, it is natural to convert between them if we take the standard interpretation of income subsidies as coming from a fixed endowment that grows geometrically to support a fixed rate of payout.

The optimal solution here directly applies the solutions for optimally allocating income and wealth subsidies shown above. Namely, our solution involves $2n$ variables corresponding to an income and wealth subsidy variable for each of the n agents.

As above, for the min-max objective, we can proceed by applying the priority algorithm using the functions, $f_i(z_i) = \psi(c_i, u_i + z_i, \beta_i, \delta_i)$ and $g_i(x_i) = \psi(c_i + x_i, u_i, \beta_i, \delta_i)$ and hence ordering the agents by their ruin probabilities. We abide by the constraint that $\sum_i kx_i + \sum_i z_i$ is at most B.

As for the min-sum objective, we define $f_i(z_i) = -\frac{\partial \gamma}{\partial z_i}$ from Equation 4 and $g_i(x_i) = -\frac{\partial \phi}{\partial x_i}$ from Equation 5. We use these f_i, g_i in our priority algorithm, thus ordering the agents by both of the partial derivatives.

Properties of Subsidies

The solutions outlined above provide a number of different priority orderings, depending on the objective function and the nature of the subsidy (e.g., income or wealth). Since these orderings impact which agents will receive subsidies first, it is natural to ask what the relationships among these priority orderings might be. For similar reasons as in the case with no initial reserve, we again set the weights w_i to be 1.

The min-max objective orders the agents by the probability of ruin. We first show that this ordering can be the reverse of the ordering given by income or wealth.

Lemma 11. *The priority ordering given by ruin probability (and likewise the wealth subsidy solution) can be the reverse of the priority ordering given by income and wealth, even when the priority ordering by income and wealth agree.*

Proof. Given an agent i, we set $u_i, c_i = 1 + i\epsilon$ and $\delta_i = 1/c_i$. The priority orderings by income and wealth agree and have ordering $(1, 2, \dots, n)$.

On the other hand, we note that the ruin probability for each agent i is given by:

$$\frac{\beta_i}{c_i \delta_i} e^{\left(\frac{\beta_i}{c_i} - \delta_i\right)u_i} = \frac{\beta_i}{e^{(1-\beta_i)}}.$$

We set $\beta_i = \frac{1}{n-i+2} - \epsilon$. Note that $\beta_i \in (0, 0.5)$. As i increases, the ruin probability above also increases. Therefore, the priority ordering given by the ruin probability is the exact reverse of the orderings given by the agents' income and wealth.

To note a similar result for the wealth subsidy, we first recall that each agent's value for Equation 4 is given by

$$\left(\frac{\beta_i}{c_i} - \delta_i\right) \frac{\beta_i}{c_i \delta_i} e^{\left(\frac{\beta_i}{c_i} - \delta_i\right)u_i} = \left(\frac{\beta_i^2 - \beta_i}{c_i}\right) e^{\beta_i - 1}$$

This equation has root $\beta_i = \frac{\sqrt{5}-1}{2}$. Thus, setting the β_i as above, we note that this value decreases as i increases, giving us a priority ordering that coincides with the ordering given by the ruin probability. $\square$

Using the case when the agents' wealth is zero, we have shown that the priority ordering by ruin probability can be the reverse of the ordering given by the income subsidy solution. We can show a similar inversion with the wealth subsidy and the ruin probability using the case where shocks are drawn from an exponential distribution. In fact, we find that a stronger result holds.

Lemma 12. *The priority ordering given by the ruin probability can be the reverse of the priority ordering given by the optimal solution for income subsidy and wealth subsidy, even when the latter two coincide.*

Proof. We let $\delta_i, u_i = 1$. For notational convenience, we set $r_i = \beta_i / c_i$. The ruin probability is given by

$$r_i e^{r-1}.$$

Equation 4 yields

$$(r_i - 1) r_i e^{r_i - 1}.$$

Note again that this has root $r_i = \frac{\sqrt{5}-1}{2}$.

On the other hand, Equation 5 yields

$$-\frac{\beta_i e^{\frac{\beta_i}{c_i^2} - 1}}{c_i^2} - \frac{\beta_i e^{\frac{\beta_i}{c_i^2} - 1}}{c_i^3}.$$

We set $\beta_i = \left(\frac{\sqrt{5}-1}{2} + \epsilon i\right) c_i$ and let $c_i = 1 - \frac{1}{i+1} - \epsilon$. Therefore, $r_i = \frac{\sqrt{5}-1}{2} + \epsilon i$. We note that the ruin probabilities are increasing in i.

On the other hand Equations 4 and Equation 5, used to determine the priority ordering for wealth and income subsidies are both increasing in i, giving us a priority ordering that is the reverse of that obtained by using the ruin probabilities. □

A distinct question, and one whose answer is not a priori clear, is whether the priority orderings for income and wealth subsidies must be the same. We see that they are, in fact, not, indicating that which agents we prioritize depends on the kind of subsidy that is being allocated by the planner.

Lemma 13. *The priority ordering given by the optimal solution for income subsidy can be the reverse of the priority ordering given by the optimal solution for the wealth subsidy.*

Proof. For each agent i, we let $c_i, u_i, \delta_i = 1$. Evaluating Equation 4, we obtain

$$\frac{(\beta_i - 1)\beta_i e^{\beta_i}}{e}.$$

And, evaluating Equation 5, we obtain

$$-\beta_i e^{\beta_i - 1} - \beta_i^2 e^{\beta_i - 1}.$$

Note that in the region $\beta_i \in [0.7, 0.8]$, the first value is increasing while the second value is decreasing. Therefore, if we set $\beta_i = 0.7 + \frac{0.1i}{n}$, for each agent i, the optimal solution for wealth subsidy would prioritize agents with smaller indices i and the optimal solution for income subsidy would prioritize agents with higher indices. □

These contrasting prioritizations surface a natural question about whether the quantities we are considering — ruin probability, optimal income subsidy value, optimal wealth subsidy value — are monotone in an agent's parameters. It is easy to note that the ruin probability exhibits monotonicity in all the parameters c_i, u_i, β_i, and δ_i. That is, holding all other variables constant, if an agent's income or wealth is lower, or they experience more frequent or larger shocks, then the ruin probability increases. However, we find that this is not the case for the optimal income and wealth subsidy values, which are only monotone in (complementary) subsets of the parameters. Table 1 contains a summary of these results.

Type	c	u	β	δ
Wealth subsidy	no	yes	no	no
Income subsidy	yes	no	yes	yes
Ruin probability	yes	yes	yes	yes

Table 1: Summary of monotonicity results for Equations 4 and 5 as well as ruin probabilities when shocks are drawn from an exponential distribution.

We can establish the monotonicity results as follows. For Equation 4 evaluated at $z_i = 0$, we get

$$\left(\frac{\beta_i}{c_i} - \delta_i\right) \psi(c_i, u_i, \beta_i, \delta_i).$$

Since u_i only shows up in ruin probability, which is monotone in wealth, this above expression is also monotone in wealth. Likewise, Equation 5 evaluated at $x_i = 0$ gives us:

$$-\left(\frac{w_i \beta_i \, \mathrm{e}^{-u_i\left(\delta_i - \frac{\beta_i}{c_i}\right)}}{\delta_i {c_i}^2} + \frac{\beta_i^2 \, u_i \, \mathrm{e}^{-u_i\left(\delta_i - \frac{\beta_i}{c_i}\right)}}{\delta_i {c_i}^3}\right).$$

It is easy to see monotonicity by β_i, c_i, and δ_i.

We show non-monotonicity results using the following example with three agents:

Example 14 (Wealth Subsidy)**.** We first show that Equation 4 is not monotone in c_i, β_i, δ_i. Suppose that $w_i, c_i, u_i, \delta_i = 1$ and $\beta_i = 0.5$ for each agent. Then, the value for Equation 4 is -0.1516. To see non-monotonicity, we change the values for c_i, β_i, δ_i as follows and evaluate Equation 4:

c: Set $c = (0.6, 1, 2)$, which evaluate to

$$(-0.1176, -0.1516, -0.0886).$$

β: Set $\beta = (0.1, 0.5, 0.9)$, which evaluate to

$$(-0.037, -0.152, -0.081).$$

δ: Set $\delta = (0.6, 1, 2)$, which evaluate to

$$(-0.0754, -0.1516, -0.0837).$$

Each of these induce a priority ordering $(2, 1, 3)$ for the wealth subsidy instead of the ordering $(1, 2, 3)$ that would be implied by monotonicity.

We can show non-monotonicity in u_i for Equation 5 using the following example with three agents.

Example 15 (Income Subsidy)**.** Suppose we have agents where $w_i, c_i, \delta_i = 1$ and $\beta = 0.7$. Set their wealth to be $(0.1, 2, 5)$. Equation 5 evaluates to $(-0.7269, -0.9220, -0.7029)$, giving us the claimed non-monotonicity result.

D Non-zero Wealth: General Distributions

Proposition 7. Let MIN-SUM take as input explicit descriptions of n functions $f_i(\cdot)$, and output $\min_{\vec{x}, x_i \geq 0 \forall i, \sum_i x_i \leq B} \{\sum_i f_i(x_i)\}$. Then, MIN-SUM is (weakly) NP-hard, even when each $f_i(\cdot)$ takes the form of $\min\{1, g_i(\cdot)\}$, where $g_i(\cdot)$ is convex.[3]

Proof. We reduce from SUBSET-SUM. Given an instance of SUBSET-SUM with $A_1, \ldots, A_n$ for which we want to know whether there is a subset that sums to exactly C, we do the following:

[3]Observe that without the minimum with 1, this is just convex minimization, which can be solved in polynomial time.

- Let $A \geq \max_i A_i$.
- Define $g_i(x_i) := \max\{1 - A_i/A, 1 + (A_i - 2x_i)/A\}$. Observe that $g_i(\cdot)$ is the maximum of linear functions and therefore convex. Now define $f_i(x_i) := \min\{1, g_i(x_i)\}$. Note that $f_i(0) = 1$.
- Input $f_1(\cdot), \ldots, f_n(\cdot), C$ to MIN-SUM. If there is a solution with value $n - C/A$, answer "yes" to SUBSET-SUM. Otherwise, answer "no."

The above reduction clearly runs in polynomial time. Therefore, it only remains to show whether it is correct. First, observe that the derivative of $f_i(\cdot)$ is either 0 or $-2/A$ everywhere. In particular, $f_i'(x_i) = 0$ on $[0, A_i/2]$, $f_i'(x_i) = -2/A$ on $[A_i/2, A_i]$, and $f_i'(x_i) = 0$ on $[A_i, \infty)$.

Therefore, if we look at a candidate solution $\vec{x}$, we necessarily have $f_i(x_i) \geq 1 - x_i/A$. This is simply because the average derivative on $[0, x_i]$ is at least $-1/A$ for all x_i. The inequality is tight only at $x_i = A_i$. From this, we conclude that if there is a subset such that $\sum_{i \in S} A_i = C$, we could consider setting $x_i = A_i$ for all $i \in S$ and $x_i = 0$ otherwise. This would yield a solution with value $n - C/A$. This is because all $i \in S$ will have $f_i(x_i) = 1 - A_i/A$, and so our total solution will have value $n - C/A$.

Similarly, if there exists a solution with value $n - C/A$, this implies that the gain in value for every unit of budget is exactly $-1/A$ (because no budget contributes better than $-1/A$, and our total budget is B). For each agent, this implies that either they consume no budget ($x_i = 0$) or they consume exactly A_i of budget. Because the entire budget is consumed at bang-for-buck of $-1/A$, if S denotes the agents who receive non-zero subsidy we must have $\sum_{i \in S} A_i = C$. Therefore, the reduction is correct. □

Lemma 10. If all $\psi_{u_i,\beta_i,F_i}(\cdot)$ can be computed in $O(1)$ operations. Then if X denotes the optimal solution to the min-max objective with budget B, a solution with min-max value $X + \delta$ using budget $B + \delta$ can be found in time $\mathrm{poly}(n, \log(B/\delta))$. If Y denotes the optimal solution to the min-sum objective with budget B, a solution with min-sum value $(1 + \varepsilon) \cdot Y + \delta$ using budget $B + \delta$ can be found in time $\mathrm{poly}(n, 1/\varepsilon, \log(B/\delta))$.

Proof. Similar to Corollary 9, the proof follows by observing that a $y \leq \psi^{-1}_{u_i,\beta_i,F_i}(x) \leq y + \delta/n$ can be found in $O(\ln(nB/\delta))$ operations. Using $y + \delta/n$ as $\psi^{-1}_{u_i,\beta_i,F_i}(x)$ results in no additional error in ruin probabilities, but does cost at most an additional δ/n in budget consumed. Chasing through the rest of the previous proofs, the same guarantees are achieved, but after consuming an additional δ budget. □

E Societal Implications

There are a number of societal considerations inherent in the problem we are studying. One such consideration is connected to the assumption that shocks can be observed or their distributions can be known in some instances. As discussed in the introduction, obtaining this information can be intrusive into the lives of vulnerable individuals. While some information, such as interactions with the criminal justice system, may already be data that a planner — such as a government assistance program — already has access to, others — such as the dissolution of a romantic or other close personal relationship — may represent private and sensitive information. There is a rich line of work exploring the class differentials in privacy loss and privacy violations, especially for data gathered as part of government assistance programs (Gilman 2011; Marwick, Fontaine, and boyd 2017; Peppet 2011). This suggests a number of open questions, including how to allocate subsidies effectively when we only have noisy data about agents — including cases where the noise comes from the planner deliberately limiting their data collection on shocks.

A separate issue is concerned with setting the objective function. In this work, we study the min-max and min-sum objectives, which aim to maximize the welfare of the agent who is most susceptible to experiencing ruin and minimizing the expected number of agents that may experience ruin. Both of these are well-studied objectives in the literature and also correspond to objectives that are motivated in practice. For instance, in the motivating literature related to eviction, a planner might aim to minimize the maximum likelihood that any family will get evicted or the planner may wish to minimize the expected number of families who may experience eviction. Social programs that work by identifying and targeting families who may benefit most from such programs are common in various domains, including low-income housing assistance programs and poverty-reduction efforts more generally.

At the same time, social policy may also provide such provisions to all families. There is rich work across the social sciences and public policy exploring the role of targeting in poverty reduction ranging from broader survey (Akerlof 1978; Elster 1992; Mirrlees 1971), to the role of self-targeting and other mechanisms that may drive the process (Alatas et al. 2016; Nichols and Zeckhauser 1982), to work focused on specific domains such as medicine (Persad, Wertheimer, and Emanuel 2009). Note that in some domains, there may be a choice between policy that takes a "targeting" versus "universalist" approach, where programs would build a universal floor for all families (Mkandawire 2005; Skocpol 1991).

Recent work including by Eubanks has highlighted the complex interaction between public services and algorithmic decision-making tools as well as challenges that arise in implementing such solutions in real-world settings (Eubanks 2018a; 2018b). Prior to this work, researchers and practitioners in economics and computation have studied the societal and ethical considerations that emerge when allocating scarce societal resources and the role of repugnant markets (Calabresi and Bobbitt 1984; Elster 1992; Roth 2007; 2008).