

Subsidy Allocations in the Presence of Income Shocks

Rediet Abebe,^{1,2} Jon Kleinberg,² S. Matthew Weinberg³

¹Harvard University

²Cornell University

³Princeton University

Abstract

Poverty and economic hardship are understood to be highly complex and dynamic phenomena. Due to the multi-faceted nature of welfare, assistance programs targeted at alleviating hardship can face challenges, as they often rely on simpler welfare measurements, such as income or wealth, that fail to capture to full complexity of each family’s state. Here, we explore one important dimension – *susceptibility to income shocks*. We introduce a model of welfare that incorporates income, wealth, and income shocks and 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 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 designed to address the needs of vulnerable individuals and families (Anand and Sen 1997; Atkinson 2003). A crucial challenge here 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 this 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; 2016), which surveys approximately 2300 families in New York City, documenting the intricate associations between their circumstances and levels of hardship. As with other studies in this area, the Poverty Tracker study is in part based on the premise

that we need to broaden our frameworks for quantifying economic well-being; the reports state 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).

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 state (Giesbert and Schindler 2012; Kochar 1995; O’Flaherty 2009). Such shocks can be a result of an unexpected expense (e.g., a parking ticket or medical expense), a delayed paycheck or unexpected loss of a job, loss of a romantic 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 (Desmond 2012; 2016; Dinkelman, Lam, and Leibbrandt 2008; Kochar 1999; O’Flaherty 2009). 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 significance 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 assistance 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? There is a long line of work in the study of household consumptions and public economics that considers agent-level behavior in response to subsidies and more broadly understanding consumption dynamics (Aiyagari 1994; Golosov et al. 2006; Kocherlakota 2004) as well as optimal taxation theory (Eaton and Rosen 1980a; 1980b; 1980c; Farhi and Werning 2013; Saez and Stantcheva 2018; Varian 1980). The work presented here is informed by these areas. We however take an algorithmic approach and study the inherent stochastic optimization problems 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.

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 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 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, our algorithms 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 type of subsidy. We close with a discussion on open questions as well as societal considerations related to optimization-informed 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 vulner-

ability. 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}, \dots 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

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

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

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

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 full paper as well as (Abebe et al. 2020).

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. In this setting, the ruin probability will then be 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. This mathematical assumption is thus 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 $\psi(c_i, 0, \beta_i, F_i) = \psi(c_i, 0, \beta_i, F_i')$ (Asmussen and Albrecher 2010). 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

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

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 (Alaëi 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 full paper.

Our problem is as follows:

(*) *We have functions f_1, \dots, 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, \dots, 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 $x^* = (x_1^*, \dots, x_n^*)$ 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 $f_1(0) \geq f_2(0) \geq \dots \geq f_n(0)$; 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 full paper. 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 $x^* = (x_1^*, \dots, x_n^*)$ 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 be denoted by a function ϕ where

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

Using Equation 1, we can restate this to be

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

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

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

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 $x^* = (x_1^*, \dots, x_n^*)$ 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. *For any pair of orderings given by (I) income, (II) ruin probability, or (III) our solution given by Equation 2, there exist instances on which they are reverses of one another.*

We present the proof in the full paper. On the other hand, we find that there are some dependencies among the priority orderings, including the following result.

Proposition 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 full paper for details.

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 full paper contains a discussion of results showing a range of counter-intuitive observations. We note, in particular, that wealth subsidy for the min-sum objective is only monotone in u_i and not in c_i, β_i , and δ_i . Income subsidy for the min-sum objective is monotone in the latter set and not u_i . For the min-max objective, both subsidies are monotone in each of the four 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 in this general setting. 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 $\psi_{u,\beta,F}^{-1}(\cdot)$ 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 $\psi_{u_i, \beta_i, F_i}^{-1}(p)$ for all i , and update $x_i := \max\{0, \psi_{u_i, \beta_i, F_i}^{-1}(p) - c_i\}$. 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. \square

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 non-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 $\psi_{u_i, \beta_i, F_i}^{-1}(\cdot)$ (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 $(1 + \varepsilon)\text{OPT}$, as desired. \square

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 as-

sumptions, and also discuss relaxations of the assumption that $\psi_{u, \beta, F}^{-1}(\cdot)$ 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.

Observe that without the minimum with 1, this is just convex minimization, which can be solved in polynomial time. 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 $\psi_{u_i, \beta_i, F_i}^{-1}(\cdot)$. See full paper for proofs.

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 . \square

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 $\text{poly}(n, 1/\varepsilon, \log(LB/\delta))$.

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 $\psi_{u_i, \beta_i, F_i}^{-1}(\cdot)$) 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:

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 $\text{poly}(n, \log(B/\delta))$. 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 $\text{poly}(n, 1/\varepsilon, \log(B/\delta))$.

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 2003; 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. We study susceptibility to income shocks as one possible dimension. 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.

There is a line of modeling work in economics that studies household consumption dynamics (Aiyagari 1994; Golosov et al. 2006; Kocherlakota 2004) as well as optimal taxation theory, including in the presence of uncertainty as well as optimal taxation theory (Eaton and Rosen 1980a; 1980b; 1980c; Farhi and Werning 2013; Saez and Stantcheva 2018; Varian 1980). Informed by this work, we take an algorithmic approach here and study the stochastic optimization problem inherent in this setting. 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, this is the first work to look at an optimization problem of this nature inspired by household consumption and public economics.

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 (Abebe and Goldner 2018a; 2018b). A recent study considers 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 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 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 SNAP 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 leave open directions for further work. 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. See full paper for further open directions and discussions about societal implications.

Acknowledgements

RA was supported a Facebook fellowship, a Simons Investigator Grant, and a MacArthur Foundation Grant. JK was supported in part by a Simons Investigator Award, a Vannevar Bush Faculty Fellowship, a MURI grant, and a MacArthur Foundation grant. SMW was supported by NSF-CCF grant #1717899. The authors thank members of the Cornell AI, Policy, and Practice Group as well as the Center for the Study of Inequality.

References

- Abebe, R., and Goldner, K. 2018a. Mechanism design for social good. *AI Matters* 4(3):27–34.
- Abebe, R., and Goldner, K. 2018b. A report on the workshop on mechanism design for social good. *ACM SIGecom Exchanges* 16(2):2–11.
- Abebe, R.; Barocas, S.; Kleinberg, J.; Levy, K.; Raghavan, M.; and Robinson, D. 2020. Roles for computing in social change. In *Conference on Fairness, Accountability and Transparency*.
- Aiyagari, S. R. 1994. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics* 109(3):659–684.
- 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*.
- Anand, S., and Sen, A. 1997. Concepts of human development and poverty! a multidimensional perspective. *United Nations Development Programme, Poverty and human development: Human development papers* 1–20.
- 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.
- 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.
- Desmond, M. 2012. Eviction and the reproduction of urban poverty. *American Journal of Sociology* 118(1):88–133.
- Desmond, M. 2016. *Evicted: Poverty and profit in the American city*. Broadway Books.
- Dinkelman, T.; Lam, D.; and Leibbrandt, M. 2008. Linking poverty and income shocks to risky sexual behaviour: evidence from a panel study of young adults in cape town. *South African Journal of Economics* 76:S52–S74.
- Eaton, J., and Rosen, H. S. 1980a. Labor supply, uncertainty, and efficient taxation. *Journal of Public Economics* 14(3):365–374.
- Eaton, J., and Rosen, H. S. 1980b. Optimal redistributive taxation and uncertainty. *The Quarterly Journal of Economics* 95(2):357–364.
- Eaton, J., and Rosen, H. S. 1980c. Taxation, human capital, and uncertainty. *The American Economic Review* 70(4):705–715.
- Farhi, E., and Werning, I. 2013. Insurance and taxation over the life cycle. *Review of Economic Studies* 80(2):596–635.
- Findeisen, S., and Sachs, D. 2016. Optimal need-based financial aid. Technical report, Mimeo, University of Mannheim/European University Institute.
- Giesbert, L., and Schindler, K. 2012. Assets, shocks, and poverty traps in rural mozambique. *World Development* 40(8):1594–1609.
- Golosov, M.; Tsyvinski, A.; Werning, I.; Diamond, P.; and Judd, K. L. 2006. New dynamic public finance: A user’s guide [with comments and discussion]. *NBER macroeconomics annual* 21:317–387.
- 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.
- Kochar, A. 1995. Explaining household vulnerability to idiosyncratic income shocks. *The American Economic Review* 85(2):159–164.
- Kochar, A. 1999. Smoothing consumption by smoothing income: hours-of-work responses to idiosyncratic agricultural shocks in rural india. *Review of Economics and Statistics* 81(1):50–61.
- Kocherlakota, N. R. 2004. Wedges and taxes. *American Economic Review* 94(2):109–113.
- Kube, A. R.; Das, S.; and Fowler, P. J. 2018. Allocating interventions based on counterfactual predictions: A case study on homelessness services.
- O’Flaherty, B. 2009. What shocks precipitate homelessness. *Unpublished manuscript*. New York: Columbia University.
- Saez, E., and Stantcheva, S. 2018. A simpler theory of optimal capital taxation. *Journal of Public Economics* 162:120–142.
- 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.
- Varian, H. R. 1980. Redistributive taxation as social insurance. *Journal of public Economics* 14(1):49–68.
- 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.
- Wimer, C.; Collyer, S.; Garfinkel, I.; Maury, M.; Neckerman, K.; Teitler, J.; Waldfoegel, J.; and Wutchiett, D. 2016. Poverty tracker – monitoring poverty and well-being in NYC: Report 4. Technical report.