Optimal Redistribution Through Public Provision of Private Goods[*]

Zi Yang Kang[†]

This version: August 2021

First version: November 2020

Abstract

How does a private market influence the optimal design of a public program? In this paper, I study a designer who has preferences over how a public option and a private good are allocated. However, she can design only the public option. Her design affects the distribution of consumers who purchase the private good—and hence equilibrium outcomes. I characterize the optimal mechanism and show how it can be computed explicitly. I derive comparative statics on the value of the public option and show that the optimal mechanism generally rations the public option. Finally, I examine implications on the optimal design when the designer can intervene in the private market or introduce an individual mandate.

JEL classification: C61, D47, D61, D63, D82

Keywords: public option, redistribution, rationing, public provision, mechanism design

---

1 Introduction

Many public programs operate alongside private markets, allowing consumers to choose between the public option and the private good. For example, eligible individuals may either rent an affordable housing unit through a public housing program or lease a private apartment at the market price. Similarly, even though public health care and education are available in many countries, individuals may instead choose to seek private care or attend a private school. The outcome of any public program therefore depends on equilibrium allocations of both the public option and the private good; yet most theoretical models, which adopt a partial equilibrium analysis, do not account for equilibrium effects of the design on the private market.

In this paper, I study how a designer optimally provides a public option in equilibrium with a competitive private market. The designer supplies the public option at a constant marginal cost, as long as the total cost of doing so, net of any revenue, does not exceed her budget. Consumers have heterogeneous preferences over the public option, the private good, and money, but value only a single unit of the good: they want either an affordable housing unit or a private apartment, but not both. Consumer utility is linear over the good and money. All consumers can—but are not required to—purchase the public option. After allocations for the public option are realized, consumers who receive the public option leave the market. The remaining consumers proceed to the private market, where the competitive equilibrium is realized.

The designer maximizes total utility, which is equivalent to the sum of weighted consumer surplus and producer surplus, with weights on consumers that are functions of their rates of substitution between the public option and the private good, and between the private good and money. While her objective depends on the joint allocation of the public option and the private good, the designer is only able to allocate the public option. Although the designer is unable to intervene directly in the private market, her choice of mechanism does influence the equilibrium price that will prevail in the private market—and hence the outside option of participants in the mechanism.

To see why the consumers' outside option depends on the designer's choice of mechanism, suppose that the designer runs a lottery that allocates the public option for free to every consumer who participates with equal probability. In the absence of a private market, all consumers participate in the lottery, which allows the designer to easily evaluate the outcome of such a design. However, the logic is not as straightforward once a private market is introduced. Not all consumers participate in the lottery: depending on the price of the private good, some consumers might have a negative option value for the public option—even when it is free. The

price of the private good varies with residual demand after the public option has been allocated, which in turn depends on which consumers participate in the lottery. Thus, when a private market is present, the designer must solve a fixed-point problem in order to identify the mapping from mechanisms to outcomes.

While solving for the optimal mechanism might appear complicated in light of these equilibrium effects, the problem is, perhaps surprisingly, tractable. The main result of this paper shows that optimal provision of the public option is achieved through a small number of rationing options. Specifically, the designer's optimal menu consists of at most three prices. Consumers who pay the high price are allocated the public option with probability 1. Consumers can also choose to pay either the medium or low price, in which case they are rationed: they are allocated the public option with probability less than 1.

The key observation that simplifies the analysis is that the price in the private market and the allocation rule for the public option jointly determine consumer decisions. This observation allows the design problem to be decomposed into two stages: the designer first chooses the price to effect in the private market, and then chooses an allocation rule for the public option—subject to the constraint that it effects the chosen price. Such a decomposition is useful because it allows us to focus on the designer's second-stage problem, which can then be formulated as a mechanism design problem with two constraints: the designer's budget constraint and the constraint that the mechanism effects the designer's chosen price in the private market. As I show in Section 3, these constraints are affine in the allocation rule, so the essential insight of Myerson (1981) applies: the design problem is an infinite-dimensional linear program, and thus an optimal allocation rule exists at an extreme point of the set of all feasible allocation rules. Because of the additional constraints in the design problem, these extreme points are allocations that are implementable not by a single price, but rather by at most three prices.

While this analysis provides a transparent characterization of the optimal mechanism, it leaves unanswered several important questions from a practical perspective. A central question to many a policy debate is the value of the public option: how much total utility is gained through the optimal public provision mechanism? In Section 4, I answer this by exploiting convex duality. Strong duality holds for the design problem: the value of the design problem is equal to the value of its dual problem, which assigns to each constraint a shadow price. The value of the dual problem can be found by computing the optimal shadow price of each constraint. Because the dual problem is convex and has only two shadow prices, the resulting computation problem can be solved quickly. The value of the dual problem gives the total weighted surplus attained by the optimal mechanism, from which the value of the public option can be deduced.

As one might expect, the value of the public option is decreasing with the marginal cost of supplying the public option and increasing with the weight that the designer places on consumer surplus relative to producer surplus. Because the weight on each consumer is equal to the designer's expectation of their value for money, one interpretation of the latter comparative static is that the value of the public option decreases with consumer "wealth."

A subtler question is whether the value of the public option increases with consumer "inequality," appropriately defined. The answer is generally no: incentive constraints imply that consumers who gain the most utility from the public option are neither the poorest consumers (who have the lowest value for the good, public or private) nor the wealthiest (who have the highest outside option available), but rather "middle-class" consumers. Consequently, the value of the public option need not increase with inequality if most of this value goes to middle-class consumers rather than the poorest consumers.

While the dual approach helps determine the value of the public option, it is also of practical importance to explicitly compute the optimal mechanism—rather than to simply characterize its structure. In Section 5, given the distribution of consumer preferences and supply-side parameters, I show how to make this computation using the mathematical concept of concavification familiar from other contexts (Aumann and Maschler, 1995 and Kamenica and Gentzkow, 2011).

The use of concavification here, however, differs from how it has been used in the literature. I begin by *augmenting* the design problem to include a term that depends on the square of the allocation function. Intuitively, the augmented problem models a designer who penalizes variance in allocation probability across consumers, and reduces to the original problem in the limit where the penalty vanishes.

Paradoxically, the augmented quadratic program is easier to solve than the linear one. The augmented quadratic program can be viewed as a *regression* problem that minimizes the squared difference between the designer's chosen allocation and a "target" allocation; hence its solution reduces to projecting the designer's target allocation onto the set of all feasible allocations. This is equivalent to concavifying the target allocation (to make it monotone), and then truncating it on the unit interval (to make it a probability). By solving a sequence of augmented problems this way, the optimal mechanism for the original design problem can be computed as the pointwise limit of the solutions to the augmented problems.

So far, the analysis has prioritized generality: it allows for multiple dimensions of heterogeneity in consumer preferences. In Section 6, I take the opposite approach of imposing restrictions on consumer preferences in order to obtain economic insights on properties of the optimal mechanism. Under some regularity conditions, I show that the optimal mechanism generally requires the use

of rationing in public provision. There are two effects: a redistribution effect, which always acts in the direction of uniform randomization across all consumers, and a revenue effect arising from the designer's budget constraint, which could act in different directions depending on what price the designer chooses to effect in the private market. Nevertheless, there is an open interval of prices in the private market for which some rationing is optimal, where each price can be rationalized by an appropriate weighted producer surplus function.

In theory, the need for rationing might be reduced when the designer is able to tax or subsidize the private market, or when the designer can pass an individual mandate that requires consumers to buy either the public option or the private good. However, there are nuances to this intuition. Under general conditions, the ability to partially intervene in the private market reduces the maximum number of prices required by the optimal mechanism from three to two; however, consumers are *always* rationed, even if they pay the higher price. An individual mandate means that the revenue effect described above always acts in the direction of posting a single price for the public option. While rationing might still be optimal, I show that the scope of rationing is greatly reduced under an individual mandate.

Finally, I discuss policy implications in Section 7 and show how my results yield intuitions for designing real-world markets. Section 8 concludes.

Related literature

This paper is most closely related to the work by Dworczak Ⓡ Kominers Ⓡ Akbarpour (2021) and Akbarpour Ⓡ Dworczak Ⓡ Kominers (2020). Dworczak Ⓡ al. study the optimal mechanism for redistribution in a two-sided market for an indivisible good between buyers and sellers, subject to budget-balance and market-clearing constraints. The optimal mechanism in their setting sets at most three prices in total for both sides of the market. Akbarpour Ⓡ al. study a similar setting with goods of heterogeneous quality and publicly observed consumer characteristics. They derive the optimal mechanism and analyze its implications for the optimal trade-off between efficiency and redistribution.

Like these papers, I study the question of optimal redistribution in a market design setting where the designer accounts for heterogeneous Pareto weights across consumers. Unlike these papers, however, I model the equilibrium effect of a private market on the design. Each consumer in my model faces an outside option—equal to their payoff in the private market—that is determined endogenously by the chosen design. These endogenous outside options have important economic consequences: they change the incentives that consumers face by reordering consumer preferences.

This yields an additional force for rationing in the optimal mechanism that is not present in the analyses of Dworczak Ⓡ al. and Akbarpour Ⓡ al.

There is a substantial body of work that studies optimal allocation problems when the design objective differs from either efficiency or revenue, dating back to Weitzman (1977). Condorelli (2013) studies the use of market and non-market mechanisms when the designer distinguishes between consumer value and willingness to pay; such a distinction is possible when consumers have budget constraints, as pointed out by Che, Gale, and Kim (2013b). My paper—as well as Dworczak Ⓡ al.'s and Akbarpour Ⓡ al.'s—also borrows from the public finance literature, including contributions by Diamond and Mirrlees (1971), Atkinson and Stiglitz (1976) and Saez and Stantcheva (2016), in the use of nonuniform Pareto weights in the design objective.

The redistributive potential of public provision is a key premise of this paper, and is studied in work by Blackorby and Donaldson (1988), Besley and Coate (1991), Munro (1991), Epple and Romano (1996), and Gahvari and Mattos (2007), among others; this literature is surveyed by Currie and Gahvari (2008). This paper complements these prior contributions by taking a market-design approach to determine the optimal mechanism for providing the public option.

There has also been some related work on partial mechanism design, or "mechanism design with a competitive fringe," pioneered by Philippon and Skreta (2012), Tirole (2012), and Fuchs and Skrzypacz (2015), who study optimal interventions in markets with adverse selection. While my model shares some modeling techniques in common with these papers, the questions posed and results obtained are very different: for example, Philippon and Skreta and Tirole find that intervention in private markets cannot improve welfare, while Fuchs and Skrzypacz study if government intervention can be optimal in a dynamic setting.

Three recent contributions in partial mechanism design are Carroll and Segal (2019), Dworczak (2020), and Loertscher and Muir (2020). Among these papers, my paper is more closely related to Loertscher and Muir's, who analyze the optimal selling mechanism in the presence of a resale market. The probability of resale occurrence—as well as the potential gains from trade from resale—changes the designer's optimal mechanism. Loertscher and Muir characterize the optimal selling mechanism when the resale market is perfectly competitive, but show that other conduct assumptions on the resale market prove less tractable.

In contrast to the setting of Loertscher and Muir, tractability in my model derives from the observations that: *(i)* the price in the private market (together with the chosen allocation rule) determines the private market outcome, and *(ii)* the price in the private market can be expressed as a moment constraint of the allocation rule. These observations also explain why Loertscher and Muir encounter more tractability under some conduct assumptions on the resale market but

not others: the assumptions that they find tractable are exactly those which admit a statistic that can be expressed as a moment constraint of the allocation rule.

A growing empirical literature also examines the question of how private goods can be optimally provided by government agencies, including work by van Dijk (2019) on housing, Dinerstein, Neilson, and Otero (2020) on education, Jiménez-Hernández and Seira (2021) on milk markets, and Atal, Cuesta, González, and Otero (2021) on pharmaceutical drugs. This paper complements this line of work by showing that rationing the public option can be optimal and by providing a way to explicitly compute the optimal mechanism.

From a methodological perspective, this paper builds on techniques developed in mechanism design and information design. Following the seminal contributions of Hotelling (1931), Mussa and Rosen (1978), Myerson (1981), Maskin and Riley (1984), and Bulow and Roberts (1989), the mechanism design literature has developed a toolkit of "generalized ironing" techniques. This includes papers by Skreta (2006), Manelli and Vincent (2007), Toikka (2011), Hartline (2013), Dworczak Ⓡ Kominers Ⓡ Akbarpour (2021), and Loertscher and Muir (2020). This line of work is closely connected with the concavification techniques developed in information design, such as by Aumann and Maschler (1995), Kamenica and Gentzkow (2011), Le Treust and Tomala (2019), Doval and Skreta (2021), and Kleiner, Moldovanu, and Strack (2021).

While these techniques (particularly those by Doval and Skreta and Dworczak Ⓡ al.) apply in special cases of my model, they fail to apply in general. Dworczak Ⓡ al. use the observation that the quantity of trade can be interpreted as a "Bayes plausibility" condition, and derive the optimal mechanism by concavifying the objective function. The corresponding market-clearing constraint in my model applies only to consumers who purchase the private good, yet the designer's objective function also includes consumers who cannot afford the private good. As formalized in Sections 2 and 3, the measure defined on effective types to which the market-clearing constraint applies might not be absolutely continuous with respect to the measure defined on all effective types. Nevertheless, these techniques are useful to derive intuitions for the optimal mechanism under additional regularity conditions. I show how this can be done in Section 6.

My analysis in Sections 3, 4, and 5, however, goes beyond these standard techniques in two ways: *(i)* by characterizing the structure of the optimal mechanism without concavification, and *(ii)* by showing how the optimal mechanism can be explicitly computed. The first relies on the key observation is that concavification is not actually required for a characterization of the optimal mechanism. Instead, the simple structure of the optimal mechanism arises because design constraints can be expressed as moment constraints of the allocation rule. The design problem is therefore mathematically equivalent to a "generalized moment problem," or an infinite-dimensional

6

linear program with a finite number of moment constraints. As I show in Section 3, this can be solved using a combination of results by Bauer (1958) and Szapiel (1975), which takes advantage of the mathematical fact that a solution to a linear program is found at the extreme point of the feasible set. The second exploits the structure of infinite-dimensional *quadratic* programs, rather than linear programs, to compute the optimal mechanism. This is done in Section 5, where the optimal mechanism is obtained by solving a sequence of quadratic programs and taking the pointwise limit. Although concavification is used here, the reason for its relevance is different from its standard use in the literature: here, the optimal mechanism itself—rather than the value of the design problem—is obtained via concavification.

2 Model

There is a unit mass of consumers in a market for an indivisible good. The good is supplied by both producers as a *private good* and the designer as a *public option.*

Each consumer can hold an arbitrary amount of money, but has unit demand for the good. Consumers have heterogeneous values for the private good, the public option, and money, denoted by θ, ω, and v respectively. Throughout this paper, (θ, ω, v) is referred to as the consumer's *type.* Each consumer's type lies in the set $[\underline{\theta}, \overline{\theta}] \times [\underline{\omega}, \overline{\omega}] \times [\underline{v}, \overline{v}]$.

Consumers are risk-neutral and have linear utility. If a consumer with type (θ, ω, v) holds $X \in \{0, 1\}$ units of the private good, $Y \in \{0, 1\}$ units of the public option, and M units of money, they receive a utility of

$$\max\{\theta \cdot X, \omega \cdot Y\} + v \cdot M.$$

Because utility is scale-invariant, consumer behavior is determined by their rates of substitution between the public option and the private good, and between the private good and money. These are denoted by $\delta = \omega/\theta$ and $r = \theta/v$. The joint distribution of (r, δ) is denoted by $F(r, \delta)$, which is supported on the set $[\underline{r}, \overline{r}] \times [\underline{\delta}, \overline{\delta}]$.

There are various ways to interpret δ. Foremost, δ can be thought of as a preference parameter: the distribution of δ might capture heterogeneity among consumers in how they substitute between an affordable housing unit and a private apartment. Alternatively, δ can also be interpreted as a quality parameter: an affordable housing unit might be less luxurious than a private apartment, or living in affordable housing might carry some social stigma that impacts individual utility.

By contrast, r admits a relatively straightforward as a preference parameter: the distribution of r captures heterogeneity among consumers in how they substitute between a private apartment

7

and money. Consumers could have a high r either because they have a high value θ for the good, or because they have a low value v for money.

The Marshallian surplus of a consumer who purchases the private good at a price of p is $r - p$. The contribution of such a consumer to total (social) utility is

$$\theta - v \cdot p = (r - p) \cdot v.$$

Likewise, a consumer who purchases the public option at a price of p has a Marshallian surplus of $\delta r - p$, and the contribution of such a consumer to total utility is

$$\omega - v \cdot p = (\delta r - p) \cdot v.$$

To distinguish each consumer's contribution to total utility from their Marshallian surplus, I refer to the former as the consumer's *weighted surplus* from purchasing the private good and the public option, respectively. That is, v is the marginal utility (and hence the "weight") that society—via the designer, as described below—attaches to giving an additional unit of money to the consumer. Consumers with low v are "rich" as they have a lower marginal utility of money, and hence are assigned a lower weight. Symmetrically, consumers with high v are "poor" and are assigned a higher weight.

Summing the weighted surplus over all consumers who buy the private good at a price of p yields the total weighted consumer surplus from the private good, defined by

$$\mathbf{E}_{(\theta,\omega,v)} \left[(\theta - v \cdot p) \, \mathbf{1}_{\theta - v \cdot p \geq 0} \right] = \mathbf{E}_{(\theta,\omega,v)} \left[\left(\frac{\theta}{v} - p \right) \mathbf{1}_{\theta/v - p \geq 0} \cdot v \right]$$
$$= \mathbf{E}_{(r,\delta)} \left[(r - p) \, \mathbf{1}_{r - p \geq 0} \cdot \mathbf{E}_{(\theta,\omega,v)} \left[v \, \Big| \, \frac{\theta}{v} = r, \, \frac{\omega}{\theta} = \delta \right] \right].$$

A similar expression can be derived for consumers who buy the public option. The latter identity follows from the law of iterated expectations, and motivates the definition of the *Pareto weight*:

$$\lambda(r, \delta) \equiv \mathbf{E}_{(\theta,\omega,v)} \left[v \, \Big| \, \frac{\theta}{v} = r, \, \frac{\omega}{\theta} = \delta \right].$$

The Pareto weight $\lambda(r, \delta)$ determines the weight that society attaches to the average consumer whose rates of substitution are r and δ. As Dworczak ® Kominers ® Akbarpour (2021) argue, it is economically meaningful to think of $\mathbf{E}_\delta \left[\lambda(r, \delta) \right]$ as a decreasing function in r: society attaches a

8

higher average weight to consumers with a lower willingness to pay, perhaps because willingness to pay is correlated with wealth or socioeconomic privilege.

2.1 Private good

The private good is sold in a competitive private market. At any price p, the supply of the private good is described by the supply curve $S(p)$, which is assumed to be continuous and strictly increasing, unless explicitly stated otherwise. Supply also satisfies $S(\underline{r}) = 0$: no private good is supplied when price is equal to the lowest possible rate of substitution between the private good and money.

While I do not directly model the utility of individual producers, I assume that the *weighted producer surplus* at price p is given by the upper-semicontinuous function $V(p)$. For example, if all producers have unit value for money but face different costs, then $V(p)$ would be equal to the Marshallian producer surplus,

$$V(p) = \int_{\underline{r}}^{p} S(r) \, dr.$$

However, if producers—like consumers—have heterogeneous values for money, then $V(p)$ must account for the Pareto weights that society assigns to individual producers, analogous to the Pareto weights defined above for consumers. These Pareto weights for producers would depend on each producer's costs and value for money. The present approach of modeling weighted producer surplus as a "blackbox" function $V(p)$ allows for a unified analysis of these different cases.

2.2 Public option

A designer supplies the public option at a constant marginal cost c.[1] While not required for the main result, it is natural to think of c as being sufficiently large: for example, $c > \bar{\delta} p_{\mathrm{c}}$, where p_{c} denotes the competitive price in the private market in the absence of the public option. This would imply that the public option is more inefficient relative to the private market: the designer would not supply the public option if there was no inequality (*i.e.*, if each consumer had an identical Pareto weight).

The designer faces a budget constraint of B: the cost of providing the public option, net of any revenue that she receives from selling the public option, must not exceed B.

---

[1] At some expense of simplicity, my analysis can be extended to the case where the designer faces a continuous, increasing cost function $C(Q)$, where Q is the quantity of public option supplied.

The designer does not observe the types (θ, ω, v) of individual consumers. Instead, she knows only the distribution of types—and hence the joint distribution $F(r, \delta)$ of the rates of substitution. The designer's objective function comprises:

(i) the weighted consumer surplus that consumers receive from purchasing the public option or the private good; and

(ii) the weighted producer surplus that producers in the private market receive from selling the private good, equal to $V(p)$ at a price p.

All consumers can—but are not required to—purchase the public option. After allocations for the public option are realized, consumers who receive the public option leave the market. The remaining consumers proceed to the private market, where the competitive equilibrium is realized.

2.3 Mechanism design

The designer chooses a *mechanism* (X, T), which consists of:

(i) an *allocation function* $X : [\underline{\theta}, \overline{\theta}] \times [\underline{\omega}, \overline{\omega}] \times [\underline{v}, \overline{v}] \to [0, 1]$, so that $X(\theta, \omega, v)$ specifies the probability that a consumer with type (θ, ω, v) receives the public option; and

(ii) a *payment function* $T : [\underline{\theta}, \overline{\theta}] \times [\underline{\omega}, \overline{\omega}] \times [\underline{v}, \overline{v}] \to \mathbb{R}$, so that $T(\theta, \omega, v)$ specifies the expected payment that a consumer with type (θ, ω, v) makes to the designer.

The competitive price in the private market depends only on the distribution of consumer types who are not allocated the public option. This means that the price in the private market depends on the mechanism that the designer chooses, but only via the allocation function. Consequently, the price in the private market can be expressed as a function that maps the designer's chosen allocation function X to a price $p(X)$.

By the revelation principle, it is without loss of generality to consider only direct mechanisms under which consumers truthfully report their types. Because the market is large, individual misreports of types do *not* affect the price in the private market. Thus the price in the private market is $p(X)$ regardless of whether a consumer with type (θ, ω, v) misreports. As such, the mechanism satisfies the following incentive compatibility constraint for any (θ, ω, v) and (θ', ω', v'):

$$\omega X(\theta, \omega, v) - vT(\theta, \omega, v) + (\theta - vp(X))_+ [1 - X(\theta, \omega, v)]$$
$$\geq \omega X(\theta', \omega', v') - vT(\theta', \omega', v') + (\theta - vp(X))_+ [1 - X(\theta', \omega', v')]. \tag{IC}$$

Unlike standard mechanism design problems in which the designer can design the entire market, the constraint (IC) depends nonlinearly in the allocation function X through the price $p(X)$ in the private market.

Similarly, the individual rationality constraint will depend on the the price of the private good induced by the allocation function. Because participation in the mechanism is voluntary, the designer is restricted to give each consumer a payoff no less than what the consumer would receive in the private market. Denoting $(\cdot)_+ \equiv \max\{\cdot, 0\}$, the mechanism satisfies the following individual rationality constraint for any (θ, ω, v):

$$\omega X(\theta, \omega, v) - vT(\theta, \omega, v) + (\theta - vp(X))_+ [1 - X(\theta, \omega, v)] \geq (\theta - vp(X))_+. \tag{IR}$$

3   Optimal mechanisms

In the absence of a private market, it is straightforward to show that the optimal mechanism takes a simple form:

Proposition 1. Suppose that there is no private market: $S(p) \equiv 0$. Then there exists an optimal mechanism (X^*, T^*) that is a menu of at most 2 prices, where $\operatorname{im} X^* \subset \{0, q, 1\}$ for some $0 < q < 1$.

The key intuition underlying Proposition 1 is that, because of the designer's redistribution motive, the marginal weighted consumer surplus may not be increasing in each consumer's rate of substitution between the public option and money. If the designer's budget constraint binds, then she can generally increase total weighted surplus by randomly allocating to some consumers at a lower price.[2]

In principle, the presence of a private market might induce the designer to use a more complex mechanism: after all, the incentive compatibility (IC) and individual rationality (IR) constraints are no longer linear in allocation once a private market is introduced, as described above. However, the main result of this section shows that the optimal mechanism remains simple:

Theorem 1. There exists an optimal mechanism (X^*, T^*) that is a menu of at most 3 prices, where $\operatorname{im} X^* \subset \{0, q_1, q_2, 1\}$ for some $0 < q_1 < q_2 < 1$.

---

[2] This follows from the insights of Myerson (1981) and Samuelson (1984), and has been shown in a redistributive setting by Dworczak ® Kominers ® Akbarpour (2021).

Instead of 2 prices, Theorem 1 states that at most 3 prices are needed in the optimal mechanism when a private market is introduced. Consumers who pay the highest price are allocated the public option with probability 1, while consumers who pay less are rationed.

Showing the existence of an optimal mechanism is technical and deferred to Appendix A. In the remainder of this section, I assume that an optimal mechanism exists and prove that it requires at most 3 prices.

3.1   Conditioning on the private price

The proof of Theorem 1 requires addressing the nonlinearity of the (IC) and (IR) constraints. The key to doing so is to formulate the design problem as a two-stage problem:

1. The designer chooses a price p to effect in the private market.

2. The designer chooses a mechanism, subject to the constraint that the mechanism effects a price p in the private market.

By the revelation principle, the designer can restrict attention to incentive-compatible mechanisms in her second-stage problem. Because consumers are always able to purchase in the private market at a price of p, the value of the public option is equal to its *normalized option value* to each consumer:

$$\eta = \delta r - (r - p)_+ .$$

The relevance of the normalized option value η can be gleamed by dividing (IC) by v, which yields the following constraint for all (θ, ω, v) and (θ', ω', v'):

$$\left[\delta r - (r - p)_+\right] X(\theta, \omega, v) - T(\theta, \omega, v) \geq \left[\delta r - (r - p)_+\right] X(\theta', \omega', v') - T(\theta', \omega', v').$$

This is reminiscent of the usual incentive compatibility constraint, where $\eta = \delta r - (r - p)_+$ is interpreted as the consumer's type. For this reason, I refer to η as the consumer's *effective type*.

Let $\underline{\eta}$ and $\bar{\eta}$ respectively denote the smallest and largest values of η among all consumers, and denote the distribution of effective types by $G(\eta)$. For notational simplicity, I suppress dependence on p. The following lemma is an extension of Myerson's (1981) lemma to the present context.

Lemma 1. Suppose that the designer's mechanism is constrained to effect a price p in the private market. Then it is without loss of generality for the designer to restrict attention to mechanisms (X, T) such that:

(i) $X(\theta, \omega, v) = x(\omega/v - (\theta/v - p)_+)$ for some increasing function $x : [\underline{\eta}, \overline{\eta}] \to [0, 1]$; and

(ii) $T(\theta, \omega, v) = t(\omega/v - (\theta/v - p)_+)$ for some function $t : [\underline{\eta}, \overline{\eta}] \to \mathbb{R}$ that satisfies

$$\eta \cdot x(\eta) - t(\eta) = \underline{\eta} \cdot x(\underline{\eta}) - t(\underline{\eta}) + \int_{\underline{\eta}}^{\eta} x(s) \ \mathrm{d}s \quad \text{for any } \eta \in [\underline{\eta}, \overline{\eta}].$$

Similar observations to Lemma 1 have been made in the mechanism design literature, such as by Jehiel and Moldovanu (2001), Che, Dessein, and Kartik (2013a), and Dworczak Ⓡ Kominers Ⓡ Akbarpour (2021). The intuition is simple: given that consumer behavior depends only on their effective type, the designer cannot truthfully distinguish between two consumers who have the same effective type, even if their actual types are different.

Using Lemma 1, I simplify notation below by using the *effective mechanism* (x, t) instead of the mechanism (X, T). This gives the allocation and payment as a function of consumers' effective types instead of their actual types.

By conditioning on the private price p, the designer introduces a constraint into her design problem: her chosen mechanism must effect the price p in the private market. As the following lemma shows, this is equivalent to a *market-clearing constraint*:

Lemma 2. An effective allocation function x effects a price p in the private market if and only if the following market-clearing condition is satisfied:

$$D(p) = \int_{\underline{\eta}}^{\overline{\eta}} \phi(\eta) \cdot [1 - x(\eta)] \ \mathrm{d}G(\eta) = S(p), \tag{MC}$$

where $\phi(\eta)$ denotes the probability that a consumer with effective type η is willing to buy the private good at the price p:

$$\phi(\eta) \equiv \mathbf{P}_{(r,\delta)} \left[r > p \mid \delta r - (r - p)_+ = \eta \right].$$

13

3.2 Formulating the design problem

Abusing notation, denote by $\lambda(\eta)$ the *effective Pareto weight function*, which is the expected Pareto weight given that the consumer has an effective type η. To write down the designer's objective, consider the components of total weighted surplus:

(i) **Weighted consumer surplus.** Given the effective mechanism (x, t), each consumer with effective type η receives an expected weighted surplus of

$$[\eta \cdot x(\eta) - t(\eta)] \, \lambda(\eta) + \mathbf{E}_{(r,\delta)} \left[(r - p)_+ \, \lambda(r, \delta) \mid \eta \right].$$

(ii) **Weighted producer surplus.** This is equal to $V(p)$.

Let $\underline{U}$ denote the surplus that a consumer of type $\underline{\eta}$ receives from the public option. Lemma 1 allows the effective payment function to be expressed in terms of the effective allocation function via the envelope theorem. Thus the design problem can be written as follows:

$$\max_{x, \underline{U}} \ \Omega(p, x, \underline{U}) \equiv W(p) + \underline{\Lambda} \cdot \underline{U} + \int_{\underline{\eta}}^{\overline{\eta}} \Lambda(\eta) x(\eta) \ \mathrm{d}\eta$$

$$\text{s.t.} \begin{cases} x : [\underline{\eta}, \overline{\eta}] \to [0, 1] \text{ is increasing, } \underline{U} \geq 0, \\[2mm] \int_{\underline{\eta}}^{\overline{\eta}} \phi(\eta) \cdot [1 - x(\eta)] \ \mathrm{d}G(\eta) = S(p), \\[2mm] \int_{\underline{\eta}}^{\overline{\eta}} (\eta - c) \, x(\eta) \ \mathrm{d}G(\eta) - \int_{\underline{\eta}}^{\overline{\eta}} [1 - G(\eta)] \, x(\eta) \ \mathrm{d}\eta + B - \underline{U} \geq 0. \end{cases}$$

Here, the function $W(p)$ absorbs terms that depend only on p and not on x or $\underline{U}$:

$$W(p) \equiv V(p) + \int_{\underline{r}}^{\overline{r}} \int_{\underline{\delta}}^{\overline{\delta}} (r - p)_+ \, \lambda(r, \delta) \ \mathrm{d}F(r, \delta).$$

Moreover, the *effective survival Pareto function* $\Lambda(\eta)$ is defined by

$$\Lambda(\eta) \equiv \int_{\eta}^{\overline{\eta}} \lambda(s) \ \mathrm{d}G(s).$$

In particular, when Pareto weights on consumers are uniform and equal to 1, the effective survival Pareto function is equal to the survival probability function: $\Lambda(\eta) = 1 - G(\eta)$. Denote $\underline{\Lambda} \equiv \Lambda(\underline{\eta})$, which is independent of the price of the private good.

14

3.3   Deriving optimal mechanisms

I now apply Myerson's (1981) observation that the design problem is a linear program: the design objective and constraints are linear functionals of the effective allocation function x. To formalize this, we need the following tool from infinite-dimensional concave programming based on the combined work of Bauer (1958) and Szapiel (1975):

Theorem (constrained maximum principle)**.** Let K be a convex, compact set, and suppose that $\ell : K \to \mathbb{R}^m$ is an affine function. Let $C = \ell^{-1}(\Sigma)$ be nonempty for some convex set $\Sigma \subset \operatorname{im} \Lambda$, and suppose that $\Omega : K \to \mathbb{R}$ is a continuous and quasiconvex function. Then there exists $z^* \in C$ such that $\Omega(z^*) = \max_{z \in C} \Omega(z)$ and

$$z^* = \sum_{i=1}^{m+1} \alpha_i z_i, \quad \text{where } \alpha_1, \dots, \alpha_{m+1} \geq 0, \ \sum_{i=1}^{m+1} \alpha_i = 1 \text{ and } z_1, \dots, z_{m+1} \in \operatorname{ex} K.$$

Here, $\operatorname{ex} K$ denotes the set of extreme points of the set K.

   When K is finite-dimensional and Ω is a linear function, the constrained maximum principle simplifies to a well-known result in finite-dimensional linear programming: A linear objective defined on a convex, compact set must attain its maximum at an extreme point (*i.e.*, vertex). Carathéodory's theorem implies that any extreme point of a convex, compact set with m affine constraints is a convex combination of at most $m + 1$ extreme points of the unconstrained set.
   In the present context, define the space of implementable effective allocation functions by

$$K \equiv \left\{ x \in L^1([\underline{\eta}, \overline{\eta}]) : x \overset{\text{a.e.}}{=} \widetilde{x} \text{ for some increasing function } \widetilde{x} : [\underline{\eta}, \overline{\eta}] \to [0, 1] \right\}.$$

As I show in Appendix A, K is convex and compact. It is also well-known that the set of extreme points of K consists of the step functions (cf. Skreta, 2006 and Manelli and Vincent, 2007):

Lemma 3. The function $x \in L^1([\underline{\eta}, \overline{\eta}])$ is an extreme point of K if and only if $x \overset{\text{a.e.}}{=} \widetilde{x}$ for some increasing function $\widetilde{x} : [\underline{\eta}, \overline{\eta}] \to [0, 1]$ satisfying $\operatorname{im} \widetilde{x} \subset \{0, 1\}$.

   Lemma 3 shows that the candidate solutions to the designer's *unconstrained* problem are exactly those implementable by a single price: in the absence of the market-clearing constraint and the budget constraint, the designer sets a single price for the public option.

To complete the proof of Theorem 1, observe that the designer's objective function is linear (and hence quasiconvex) and continuous in x; so the constrained maximum principle applies. Define the linear function $\ell : K \to \mathbb{R}^2$ and the convex set $\Sigma \subset \mathbb{R}^2$ by

$$
\begin{cases}
\ell(x) \equiv \left(\int_{\underline{\eta}}^{\overline{\eta}} \phi(\eta) \cdot [1 - x(\eta)] \ \mathrm{d}G(\eta), \int_{\underline{\eta}}^{\overline{\eta}} (\eta - c) \, x(\eta) \ \mathrm{d}G(\eta) - \int_{\underline{\eta}}^{\overline{\eta}} [1 - G(\eta)] \, x(\eta) \ \mathrm{d}\eta + B - \underline{U} \right), \\
\Sigma \quad \equiv \{S(p)\} \times \mathbb{R}_+.
\end{cases}
$$

Whenever $\ell^{-1}(\Sigma)$ is nonempty, the optimal effective allocation function x^* can be written as the convex combination of at most 3 extreme points of K. By Lemma 3, this implies the existence of $0 < q_1 < q_2 < 1$ such that $\operatorname{im} x^* \subset \{0, q_1, q_2, 1\}$. This completes the proof of Theorem 1.

I conclude this section with a brief methodological discussion. The above proof of Theorem 1 differs from the concavification approach taken by Dworczak Ⓡ Kominers Ⓡ Akbarpour (2021) to prove their main result, similar to Doval and Skreta's (2021) generalization of methods by Le Treust and Tomala (2019). From a technical perspective, one might wonder if their approach might apply here. In their setting, the market-clearing constraint can be rewritten as a Bayes plausibility constraint. The designer then concavifies a Lagrangian that penalizes violations of the budget constraint via a Lagrange multiplier.

This approach is indeed possible when $\phi(\eta)$ is supported on $[\underline{\eta}, \overline{\eta}]$ and the distribution $G(\eta)$ admits a density $g(\eta)$ that is supported on $[\underline{\eta}, \overline{\eta}]$. The designer's objective can then be rewritten in the usual way:

$$
\Omega(p, x, \underline{U}) = W(p) + \underline{\Lambda} \cdot \underline{U} + \int_{\underline{\eta}}^{\overline{\eta}} \frac{\Lambda(\eta)}{g(\eta)} \cdot x(\eta) \ \mathrm{d}G(\eta).
$$

However, $\phi(\eta)$ and $G(\eta)$ are derived quantities as they depend on the designer's choice of p and the joint distribution of (r, δ). A simple example in which $G(\eta)$ does not admit a density function is when $\delta \equiv 1$ for all consumers, in which case there will generally be a probability mass of consumers with effective type $\eta = p$. Other examples where $\phi(\eta)$ might be zero in $[\underline{\eta}, \overline{\eta}]$ arise when r and δ are correlated. Thus, the concavification approach requires a more involved argument to circumvent these technical difficulties. The above proof of Theorem 1 avoids these technical difficulties: the result of Theorem 1 does not rely on $\phi(\eta)$ being nonzero or $G(\eta)$ admitting a density, but rather arises from the linear structure of the design problem.

4   Value of the public option

While Theorem 1 characterizes the optimal mechanism for allocating the public option, it leaves unanswered several important questions from a practical perspective. One such question that policymakers often consider is the value of having a public option in the first place, as opposed to in-kind redistribution via lump-sum transfers for the same budget. In this section, I answer this question by computing the increase in total weighted surplus (*i.e.*, the design objective) due to a public option that is allocated optimally.

A first attempt might try to compute the optimal mechanism implied by Theorem 1, and then determine the total weighted surplus attained by the optimal mechanism. In principle, this approach is tenable: Theorem 1 shows that the optimal mechanism can be parametrized as a five-dimensional problem (comprising three prices and two rationing probabilities). However, while the resulting problem is linear in the two rationing probabilities, it is nonlinear in the three prices, making explicit computation less appealing than might first appear.

An alternative approach, which I take below, exploits convex duality. The proof of Theorem 1 shows how the designer's problem can be formulated as an infinite-dimensional linear—and hence convex—program with two constraints: the market-clearing constraint and the budget constraint. The dual problem assigns to each constraint a shadow price; the value of the public option can thus be obtained by finding the optimal shadow prices for each constraint, rather than the optimal mechanism. Because the dual problem is convex and has only two shadow prices, the resulting computation problem is fast.

4.1   The dual problem

Let $\mu \in \mathbb{R}$ and $\alpha \geq 0$ denote the Lagrange multipliers (*i.e.*, shadow prices) associated with the market-clearing constraint and the budget constraint. The Lagrange dual function (cf. Chapter 5 of Boyd and Vandenberghe, 2004) can be written as

$$h(\mu, \alpha) \equiv \max_{x \in K, \underline{U} \geq 0} \left\{ \Omega(p, x, \underline{U}) + \mu Q(x) + \alpha J(x) \right\},$$

where

$$\begin{cases} Q(x) & \equiv \displaystyle\int_{\underline{\eta}}^{\overline{\eta}} \phi(\eta) \cdot [1 - x(\eta)] \ \mathrm{d}G(\eta) - S(p), \\[2ex] J(x) & \equiv \displaystyle\int_{\underline{\eta}}^{\overline{\eta}} (\eta - c)\, x(\eta)\, \mathrm{d}G(\eta) - \int_{\underline{\eta}}^{\overline{\eta}} [1 - G(\eta)] \ \mathrm{d}\eta + B - \underline{U}. \end{cases}$$

17

Although $h(\mu, \alpha)$ requires solving an infinite-dimensional optimization problem (over $x \in K$), this actually reduces to a *unidimensional* optimization problem: $\Omega(p, x, \underline{U})$, $Q(x)$, and $J(x)$ are linear in x, and there is no additional constraint on x. Hence, it is well-known (and can be shown by appealing to the constrained maximum principle and Lemma 3) that $h(\mu, \alpha)$ is attained at some step function x°. Thus, to compute the value of $h(\mu, \alpha)$, it suffices to maximizes over $\underline{U} \geq 0$ and the set of step functions x°, which can be parametrized by a single parameter γ: $x^\circ(\eta) = \mathbf{1}_{\eta > \gamma}$.

The dual problem is thus given by minimizing the Lagrange dual function over the Lagrange multipliers:

$$\min_{\mu \in \mathbb{R}, \, \alpha \in \mathbb{R}_+} h(\mu, \alpha).$$

Because strong duality holds, the value of the dual problem is equal to the value of the designer's objective function under the optimal mechanism. Importantly, h is a convex function, allowing the dual problem (and hence total weighted surplus with the public option) to be computed using standard convex programming algorithms.

In the absence of a public option, the competitive price p_c obtains in the private market, where

$$S(p_c) = \int_{\underline{\delta}}^{\overline{\delta}} \int_{\underline{r}}^{\overline{r}} \mathbf{1}_{r > p_c} \, \mathrm{d}F(r, \delta) = 1 - F_r(p_c).$$

Here, $F_r(r)$ is the marginal distribution of r induced by the joint distribution $F(r, \delta)$. Because there is a unit mass of consumers and a budget of B to be redistributed, each consumer receives a lump-sum transfer of B. Thus the total weighted surplus is

$$h_c \equiv V(p_c) + \int_{\underline{r}}^{\overline{r}} \int_{\underline{\delta}}^{\overline{\delta}} (r - p_c)_+ \lambda(r, \delta) \, \mathrm{d}F(r, \delta) + \underline{\Lambda} \cdot B.$$

Proposition 2. The value of the public option is given by solving the following convex program:

$$\min_{\mu \in \mathbb{R}, \, \alpha \in \mathbb{R}_+} h(\mu, \alpha) - h_c.$$

4.2 Comparative statics

While Proposition 2 is practically useful for computing the value of the public option, it also paves the way for us to answer the closely related policy question: under what conditions should we expect greater value from having a public option? To answer this question, I formalize two partial orders on any two Pareto weight functions $\lambda_1(r, \delta)$ and $\lambda_2(r, \delta)$:

1. We say that there is *uniformly more wealth* associated with $\lambda_1(r, \delta)$ than $\lambda_2(r, \delta)$ if

$$\lambda_1(r, \delta) \leq \lambda_2(r, \delta) \quad \text{for every } (r, \delta).$$

A consumer population is uniformly wealthier than another if the value of money is lower in expectation for every consumer type; hence the designer puts a lower weight on all consumers.

2. We say that $\lambda_1(r, \delta)$ is *more unequal* than $\lambda_2(r, \delta)$ if $\Lambda_1(\underline{\eta}) = \Lambda_2(\underline{\eta})$ and

$$\mathbf{E}_{(r,\delta)}\left[\lambda_1(r, \delta) \mid r \geq s\right] \leq \mathbf{E}_{(r,\delta)}\left[\lambda_2(r, \delta) \mid r \geq s\right] \quad \text{for every } s \in [\underline{r}, \overline{r}].$$

A consumer population is more unequal than another if both populations have the same average Pareto weight, but Pareto weights in the more unequal population are more disperse (*i.e.*, a "mean-preserving spread" of weights in the less unequal population).[3]

Proposition 3. The following comparative statics hold for the value of the public option:

(i) The value of the public option decreases with the marginal cost c of the public option.

(ii) If the budget B is sufficiently small, then the value of the public option decreases with uniform wealth in the consumer population.

(iii) If the effective type η is decreasing in the rate of substitution r between the private good and money, then the value of the public option increases with inequality.

The intuition behind *(i)* is clear: the value of the public option increases as its cost decreases. The intuition behind *(ii)* is straightforward as well. As consumers' wealth increases uniformly, consumers unambiguously gain less utility from the public option; however, consumers also gain less utility from lump-sum redistribution of the budget. The former effect dominates when the budget is sufficiently small.

Perhaps more surprising is *(iii)*: one might expect that the value of the public option generally increases with inequality, but this is true only with additional strong assumptions on consumer preferences. This is because consumers who benefit most from the public option are not necessarily those with the lowest rates of substitution r (*i.e.*, those with the highest Pareto weights), but rather

---

[3] This is motivated by the definition of second-order stochastic dominance. The least unequal consumer population weights each consumer equally (*i.e.*, $\lambda(r, \delta) = \underline{\Lambda}$) while the most unequal consumer population places all the weight on the consumer with the highest marginal value of money (*i.e.*, $\lambda(r, \delta) = \underline{\Lambda} \cdot \delta_{\underline{r}}$, where $\delta_{\underline{r}}$ denotes the Dirac delta distribution with a point mass at $\underline{r}$).

consumers with the highest effective types η. These two groups coincide only when η is decreasing in r at the prevailing price in the private market.

The result of Proposition 3*(iii)* relates to the "sharp elbows" effect in public discourse, which describes how the middle class might benefit disproportionately from public provision.[4] In the present context, the value of introducing a public option is mitigated by the possibility that consumers with lower Pareto weights (middle-class consumers) benefit disproportionately at the expense of consumers with higher Pareto weights (poor consumers).

In some special cases, the "sharp elbows" effect has no bite. For example, if ω is proportional to v for all consumers, then η decreases with r, satisfying the condition of Proposition 3*(iii)*. This is not true, however, if ω is positively correlated with θ, in which case η could initially increase with r and then decrease with r, as considered in Section 6.

5 Computing optimal mechanisms

In addition to determining the value of a public option, another question of practical interest to a policymaker is how the optimal mechanism can be computed. Formally, given the distribution of consumer types, the marginal cost c of producing the public option, the supply curve $S(p)$ of the private good, and the weighted producer surplus function $V(p)$, is there an explicit algorithm by which a policymaker can compute the optimal mechanism, rather than as an implicit solution to a maximization problem given in Theorem 1?

This section shows how the optimal mechanism can be computed. Throughout, I make the technical assumption that $G(\eta)$ admits a density function $g(\eta)$ supported on $[\underline{\eta}, \overline{\eta}]$: otherwise, $G(\eta)$ can always be approximated by such a distribution.

The approach I take uses the concept of concavification (cf. Aumann and Maschler, 1995 and Kamenica and Gentzkow, 2011). However, unlike its standard use in mechanism design, the value of the designer's problem is not obtained via concavification (which would again require us to implicitly infer the optimal mechanism); rather, the optimal mechanism itself is obtained.

The key idea underlying this approach is to *augment* the linear program obtained in Section 3 into a quadratic program, by including a term that depends on the square of the effective allocation

---

[4] See, for example, https://www.economist.com/britain/2015/11/12/sharper-elbows.

function. The augmented design problem is

$$\max_{\underline{U} \geq 0, \, x \in K} \left\{ W(p) + \underline{\Lambda} \cdot \underline{U} + \int_{\underline{\eta}}^{\overline{\eta}} \left[\frac{\Lambda(\eta)}{g(\eta)} \cdot x(\eta) - \varepsilon \left[x(\eta) \right]^2 \right] \, dG(\eta) \right\}$$

$$\text{s.t.} \quad \begin{cases} \displaystyle \int_{\underline{\eta}}^{\overline{\eta}} \phi(\eta) \cdot [1 - x(\eta)] \, dG(\eta) = S(p), \\[2ex] \displaystyle \int_{\underline{\eta}}^{\overline{\eta}} (\eta - c) \, x(\eta) \, dG(\eta) - \int_{\underline{\eta}}^{\overline{\eta}} [1 - G(\eta)] \, x(\eta) \, d\eta + B - \underline{U} \geq 0. \end{cases}$$

Intuitively, the augmented problem models a designer who penalizes the variance in allocation probability across different consumers. Clearly, the augmented problem is equivalent to the design problem when $\varepsilon = 0$. While it may appear that augmentation only adds more complexity to the problem, it turns out that the augmented problem can be solved elegantly, as I now describe.

Let $\mu_\varepsilon^* \in \mathbb{R}$ and $\alpha_\varepsilon^* \geq 0$ be the optimal Lagrange multipliers associated with the two constraints in the augmented problem; μ_ε^* and α_ε^* can be computed easily by solving the dual problem as in Section 4. Define the function $H_\varepsilon : [\underline{\eta}, \overline{\eta}] \to \mathbb{R}$ by

$$H_\varepsilon(\eta) \equiv \frac{1}{\varepsilon} \int_{\eta}^{\overline{\eta}} \left[\frac{\Lambda(s)}{g(s)} + \mu_\varepsilon^* \phi(\eta) + \alpha_\varepsilon^* \left[s - c - \frac{1 - G(s)}{g(s)} \right] \right] \, dG(s).$$

Denote by $\mathrm{co}\, H_\varepsilon$ the concave closure of H_ε (*i.e.*, the pointwise smallest concave function that bounds H_ε from above), and let

$$x_\varepsilon(\eta) \equiv -\frac{d}{d\eta} \, \mathrm{co}\, H_\varepsilon(\eta).$$

Because $\mathrm{co}\, H_\varepsilon$ is concave by definition, x_ε is an increasing function. The following proposition shows that a truncation of x_ε solves the augmented design problem:

Proposition 4. For any $\varepsilon > 0$, the unique optimal effective allocation function that solves the augmented design problem is $x_\varepsilon^*(\eta)$, where

$$x_\varepsilon^*(\eta) = \begin{cases} 0 & \text{if } x_\varepsilon(\eta) \leq 0, \\ 1 & \text{if } x_\varepsilon(\eta) \geq 1, \\ x_\varepsilon(\eta) & \text{otherwise.} \end{cases}$$

Proposition 4 shows how the optimal effective allocation function itself can be obtained through concavification, unlike standard concavification arguments through which the value of

the design problem is obtained. The intuition underlying Proposition 4 is very different from the geometric picture described by Aumann and Maschler (1995) and Kamenica and Gentzkow (2011). The augmented design problem can be viewed as an infinite-dimensional least-squares regression problem: the designer chooses an effective allocation function that minimizes the L^2 distance to a "target" effective allocation, $\Lambda(\eta)/\left[\varepsilon g(\eta)\right]$, subject to the market-clearing constraint, the budget constraint, and the constraint that the effective allocation function is increasing. Thus the designer accommodates the first two constraints through Lagrange multipliers, and projects the resulting Lagrangian onto the set K of increasing effective allocation functions. As I show in Appendix A, the projection operator is no other than concavifying the integral of the Lagrangian and taking the negative of its gradient.

Finally, the optimal effective allocation function of the design problem can be obtained by taking a pointwise limit, thereby yielding an explicit way to compute the optimal mechanism:

Theorem 2. As $\varepsilon \searrow 0$, x_ε^* converges pointwise to an optimal effective allocation function x^*.

6 Economic implications

While Theorem 2 shows how the optimal mechanism can be computed explicitly, some properties of the optimal mechanism can be inferred even without explicit computation. In this section, I study the implications of the optimal mechanism on public provision.

Unlike previous sections which take the most general approach possible, this section imposes stronger assumptions on consumer preferences in order to make the analysis as simple as possible. In particular, I assume that $\delta \in (0,1)$ is identical across all consumers. This might arise because there is a social stigma associated with the public option that impacts the utility of all consumers the same way, or because the good is vertically differentiated with the public option being supplied at a lower quality than the private good.[5]

As δ is identical across all consumers, notation can be simplified by omitting δ as an argument. I write the Pareto weight of a consumer as $\lambda(r)$, which is a decreasing function in the rate of substitution r between the private good and money. The distribution of r is denoted by $F(r)$, which is assumed to have a decreasing density function $f(r)$ supported on the interval $[\underline{r}, \overline{r}]$. This rules out local irregularities of $f(r)$, similar to (but stronger than) the usual increasing hazard rate

---

[5] A similar analysis can be made if the public option is supplied at a higher quality than the private good, with stark differences in results. However, there seems to be few real-life examples of the public option being provided at higher quality than the private good; hence this analysis is omitted.

or Myersonian regularity assumption. It is also satisfied by many common distributions, including the uniform distribution and the truncated exponential and Pareto (power-law) distributions.

Finally, I assume that the highest rate of substitution, $\bar{r}$, is sufficiently high: at the competitive price p_c where $1 - F(p_c) = S(p_c)$,

$$\bar{r} \geq \frac{p_c - \delta \underline{r}}{1 - \delta}.$$

Intuitively, together with the assumption that $f(r)$ is decreasing, this indicates that there is a long tail of wealthy consumers with a high value for the private good relative to money. This ensures that the wealthiest consumers always have the lowest effective types.

6.1 Rationing in public provision

The widespread use of rationing in public provision is well documented: lotteries and waitlists are common instruments for allocating public housing, while waiting times instead of prices are used to clear public health care markets. Is this optimal?

One might initially be tempted to look to Theorem 1 for an answer; but notice that it asserts only that rationing is sufficient—rather than necessary—to implement the optimal mechanism. In this subsection, however, I provide a sense in which rationing is necessary. Specifically, I show that there is an open set of prices in the private market such that the optimal mechanism allocates to a positive measure of consumers a strictly interior probability.

Before a formal proof of the result, I provide an intuition for why rationing might be optimal in public provision. To simplify the analysis, suppose for the sake of argument that the designer faces a cap on the lump-sum transfer she can make to each consumer, rather than a budget constraint. Then the only constraint on the allocation function that the designer faces is the market-clearing (MC) constraint that the mechanism effects her desired price in the private market. In this simpler setting, a stronger conclusion holds: rationing is not only optimal, but also occurs at a single price.

Proposition 5 (intuition for rationing). *Instead of a budget constraint, suppose that the designer faces a constraint on the maximum lump-sum transfer to each consumer. Then the optimal mechanism rations consumers at a single price: each consumer is allocated the public option with the same probability.*

The proof of Proposition 5 relies on two observations. First, the (MC) constraint can be written as a capacity constraint on how many units of the public option are allocated. Second, the marginal weighted total surplus gained from allocating the public option to a consumer with

effective type η is decreasing in η: even though the designer places more weight on consumers with low r, the density of consumers with middle r is higher than the combined densities of consumers with low and high r since $F(r)$ is concave.

Combining these two observations, it is easy to see that the designer optimally allocates to as many consumers with low η as possible. However, as shown in the proof of Theorem 1, any incentive-compatible mechanism must have an effective allocation function that increases with η. Thus the optimal mechanism randomizes across all consumers at a single price.

The designer's budget constraint complicates the above logic: it introduces a profit motive that acts in the opposite direction of rationing. Nevertheless, rationing remains optimal for an open set of prices[6] in the private market:

Proposition 6 (optimality of rationing). Suppose that the marginal cost c of supplying the public option is sufficiently high: $c > \delta \underline{r}$. There exists $p^\circ \leq p_{\mathrm{c}}$ such that, if the price p of the private good effected by the optimal mechanism is in a neighborhood of p°, then the optimal mechanism rations consumers.

The condition that $c > \delta \underline{r}$ is satisfied when the public option is more inefficient relative to the private market. For example, because $p_{\mathrm{c}} \geq \underline{r}$, this condition is satisfied if $c > \delta p_{\mathrm{c}}$, which in turn ensures that the designer never supplies the public option in the absence of inequality (*i.e.*, when each consumer has an identical Pareto weight).

Proposition 6 shows that the optimal mechanism generally requires rationing over a positive measure of prices—even though the profit motive introduced by the budget constraint acts in the opposite direction of rationing, and the two effects cannot be signed in general. However, the key insight to the proof of Proposition 6 is that the two effects can be signed locally, in the neighborhood of the effective type $\eta = \delta \underline{r}$, regardless of the price in the private market. There, the redistributive effect always dominates because the marginal profit changes discontinuously while the marginal weighted total surplus remains continuous. Rationing is optimal whenever the (MC) constraint implies a quantity of public option sold within a neighborhood of this discontinuity; hence there is an open set of prices over which rationing is optimal.

From a policy perspective, the results of this subsection suggest that non-market mechanisms, which may involve using non-price instruments to clear the market, merit consideration for public provision due to their potential optimality for redistribution. This is especially true if the designer

---

[6] Notice that any optimal price p^* can be rationalized by an appropriate weighted producer surplus function $V(p)$: one can simply let $V(p) = 0$ for $p < p^*$ and $V(p) = V^*$ for $p \geq p^*$, where V^* is sufficiently large.

faces no profit motive (Proposition 5), although rationing can still be optimal even with a profit motive (Proposition 6).

6.2   Intervention in the private market

So far, I have maintained that the designer can design only the public option. In many real-world settings, while the designer cannot design the private market, the designer might still be able to undertake limited interventions in the private market. In this subsection, I show how the analysis of previous sections extend when the designer is allowed to tax or subsidize the private good.

Perhaps surprisingly, the ability of the designer to intervene through a private tax or subsidy does not complicate the analysis, but rather simplifies it. To begin, observe that Theorem 1 extends to this environment:

Proposition 7 (optimal mechanism with private tax). Suppose that the designer can intervene in the private market by setting a tax τ, where $\tau < 0$ is interpreted as a subsidy for the private good. Then the optimal mechanism (X^*, T^*) for providing the public option is a menu of at most 3 prices, where $\operatorname{im} X^* \subset \{0, q_1, q_2, 1\}$ for some $0 < q_1 < q_2 < 1$.

The proof of Proposition 7 is almost identical the proof of Theorem 1, with the only difference being that the designer chooses both the tax and the price of the private good (gross of the tax), and then chooses a mechanism that effects the price at the prevailing tax. This decomposition yields a market-clearing constraint and a budget constraint that are affine in the allocation function; hence applying the constrained maximum principle yields the desired result.

Given that the designer has an additional instrument in the form of a private tax, one might wonder if the optimal mechanism simplifies: intuitively, for the purpose of redistribution, a private tax might be a substitute for rationing. This intuition is almost correct, in the sense that the optimal mechanism simplifies: one fewer price is required in the optimal menu. However, the following result shows that rationing is nonetheless required in the optimal mechanism.

Proposition 8 (rationing with private tax). Let the marginal cost c of supplying the public option be sufficiently high, so that $c > \delta\underline{r}$; and suppose that $\underline{\Lambda} \cdot S(p) \geq V'(p) \geq S(p)$ for all p. Then the optimal mechanism (X^*, T^*) for providing the public option is a menu of at most 2 prices, where $\operatorname{im} X^* \subset \{0, q_1, q_2\}$ for some $0 < q_1 < q_2 < 1$.

There are two conditions in Proposition 8. The first states that the public option is sufficiently expensive, so that the designer faces a need to raise revenue. Like the condition in Proposition 6, this condition is satisfied if the public option is more inefficient relative to the private market. The second requires that the designer weights consumer surplus more than producer surplus. For example, this condition is satisfied when $\underline{\Lambda} \geq 1$ and $V(p)$ is equal to Marshallian producer surplus, so that

$$V'(p) = S(p) \leq \underline{\Lambda} \cdot S(p).$$

Under these conditions, Proposition 8 shows that two prices—instead of three—are required in the optimal mechanism. Moreover, notice that $1 \notin \operatorname{im} X^*$: thus Proposition 8 also strengthens the result of Proposition 6 by showing that rationing is a *necessary* part of the optimal mechanism, even when the designer can tax the private good. The proof of Proposition 8 formalizes the sense in which the tax on the private good is a "substitute" for the highest tier in the optimal mechanism: namely, the designer can always increase the value of her objective function by raising the private tax if some consumers purchase the public option with certainty.

Finally, I conclude this section by showing that the designer strictly benefits from being able to tax the private good:

Proposition 9 (optimality of intervention)**.** Under the same assumptions as Proposition 8, intervention in the private market is optimal: it is never optimal for the designer to set $\tau = 0$.

Like Proposition 8, Proposition 9 is shown by contradiction: if the designer does not intervene in the private market, then she always strictly benefits by raising the tax.

In sum, the analysis of the previous sections extend to an environment where the designer can tax the private good (Proposition 7). The ability to tax the private good allows for greater redistribution (Proposition 9), but does not remove the need for rationing. By contrast, rationing remains necessary under general conditions (Proposition 8).

6.3 Individual mandates

In some instances of public provision, the designer can mandate that all consumers purchase the good, be it the public option or the private. Many countries have passed compulsory education laws, while some—including Australia, Japan, and the Netherlands—have health insurance mandates. How does an individual mandate change the optimal provision of the public option? While academic and popular discussions primarily focus on external effects and adverse selection

([Summers](), [1989]()), this subsection examines how an individual mandate changes consumer incentives via incentive compatibility and individual rationality constraints.

Under an individual mandate, consumers must purchase the private good if they are not allocated the public option. The following incentive compatibility constraint holds for any (θ, ω, v) and (θ', ω', v'):

$$\omega X(\theta, \omega, v) - vT(\theta, \omega, v) + [\theta - vp(X)] [1 - X(\theta, \omega, v)]$$
$$\geq \omega X(\theta', \omega', v') - vT(\theta', \omega', v') + [\theta - vp(X)] [1 - X(\theta', \omega', v')].$$

Moreover, the individual rationality constraint for a consumer of type (θ, ω, v) is

$$\omega X(\theta, \omega, v) - vT(\theta, \omega, v) + [\theta - vp(X)] [1 - X(\theta, \omega, v)] \geq \theta - vp(X).$$

Collectively, these incentive constraints imply that consumer effective types are now given by

$$\eta = p - (1 - \delta) r,$$

which decreases with r since $\delta < 1$ by assumption. Proposition 3 implies that:

Proposition 10 (value of public option under individual mandate)**.** Under an individual mandate, the value of the public option increases with inequality.

Proposition 10 suggests that an individual mandate helps to screen consumers more effectively: consumers' effective types are now monotone in their Pareto weight. Its intuition is identical to that of Proposition 3 and hence omitted here.

By screening consumers more effectively, one might suspect that an individual mandate is able to reduce the designer's need to ration. The following proposition confirms this intuition.

Proposition 11 (conditions for rationing under individual mandate)**.** Suppose that the following function is strictly quasiconcave for all $r \in [\underline{r}, \overline{r}]$:

$$\psi(r) \equiv \frac{1 - F(r) - \Lambda(r)/\underline{\Lambda}}{f(r)} - r.$$

Under an individual mandate, rationing is optimal only if the quantity of public option provided is sufficiently small. Moreover, it is never optimal to ration consumers at a single price.

Proposition 11 relies on a regularity condition first introduced by Dworczak ® Kominers ® Akbarpour (2021). As they argue, the concavity of $\psi(r)$ is closely related to the decreasing Pareto weights $\lambda(r)$; the two conditions are equivalent when r is uniformly distributed. Thus, the concavity $\psi(r)$ and the monotonicity of $\lambda(r)$ jointly rule out local irregularities in the density function $f(r)$. Under this condition, Proposition 11 represents an extension of Theorems 3 and 4 of Dworczak ® al. (2021) to a public provision setting. Such an extension is possible only because consumers' effective types are monotone in their willingness to pay under an individual mandate. While it does not entirely rule out rationing, Proposition 11 shows that rationing has more limited relevance under an individual mandate. This suggests that, in the absence of an individual mandate, rationing plays an important role in helping screen consumers.

In closing, it is worth cautioning against interpreting Propositions 10 and 11 as advocating for the use of individual mandates in public provision. In particular, conditional on not being allocated the public option, the poorest consumers receive lower surplus than they would without an individual mandate. The designer thus faces a tradeoff between screening consumers more effectively and lowering the surplus of consumers who are not allocated the public option. This tradeoff can be resolved by applying Proposition 2 to compute the designer's objective both with and without an individual mandate; but this tradeoff is not captured by Propositions 10 and 11.

7 Policy implications

In recent years, the public option has become a topic of keen interest for many policymakers. Even though the term "public option" has traditionally been associated with health care markets, public options have been proposed for a number of other markets, including retirement pensions, banking, childcare, broadband internet, and more (Sitaraman and Alstott, 2019). In this section, I describe how the analysis presented in this paper can be applied to some of these markets.

Housing. By interpreting consumers as renters, the public option as affordable housing units, and the private good as private apartments, the rental housing market can be viewed through the lens of the model presented in this paper.[7] Empirical studies have emphasized substantial heterogeneity in renters' incomes (and hence their price sensitivities), implying that there is

---

[7] Beyond rental housing, the framework above can also be applied to home ownership markets in places such as Hong Kong, Singapore, and South Korea, where residents can purchase ownership of houses produced under public housing schemes. In Singapore, for example, more than 80% of residents live in public housing flats, with the majority of these residents owning their flats.

significant inequality within the renter population. The quality of affordable housing units, as well as renters' perceptions of affordable housing, determines renters' rates of substitution between affordable housing units and private apartments. Renters who apply to affordable housing units are optimally rationed (Theorem 1), which might involve a combination of lotteries and waitlists.

The value of affordable housing depends on the distribution of preferences and inequality in the renter population. However, where renter populations are poorer, affordable housing might be a more effective tool than in-cash redistribution, especially with a modest budget (Proposition 3). With larger budgets, it might be feasible to provide a large quantity of high-quality affordable housing units; but doing so benefits middle-class consumers on the margin. This is the case in Amsterdam's affordable housing program, for instance, where a substantial fraction of households living in affordable housing have incomes above the median (van Dijk, 2019). Proposed policies that tax rental transactions in the private market (Diamond, McQuade, and Qian, 2019) could be an additional redistributive tool to complement the provision of affordable housing (Proposition 7).

Health care. The market for health insurance can also be evaluated using the model presented in this paper: consumers choose between a public option and a private plan. However, several theoretical abstractions that the model makes are more pronounced in this setting. For example, there could be significant market power in the private market, which is likely to increase the value of the public option (Proposition 2): the public option has the additional effect of enhancing competition in the private market by giving consumers a new option. There might also be adverse selection if there is substantial correlation between consumer preferences and health spending (Einav and Finkelstein, 2011 and Geruso and Layton, 2017).

Beyond these caveats, however, the analysis of this paper provides some intuitions for design considerations in health insurance. Public discourse on a public health insurance option in the United States has focused on the economic reasons of mitigating market power among insurers and alleviating the adverse selection problem; yet the results of this paper suggest that redistribution could be a sufficient motive for a public option in this market, independently of these issues (Proposition 5). However, a limited budget suggests that rationing might be required in general, even with an individual mandate, as this means that there is a limited quantity of the public option that can be supplied (Proposition 11). Rationing has the natural interpretations of lower coverage levels, longer wait times for care, or even a literal lottery among eligible individuals, as in the case of the 2008 Medicaid expansion in Oregon that led to the Oregon health insurance experiment.

Essential goods. Yet another market that can be approximated by the model presented in this paper is the market for essential goods, such as rice, sugar, milk, oil, and pharmaceutical drugs. Many developing countries publicly provide a public option for these essential goods in ration shops. However, the quality of the public option might be lower than that of private goods (Jiménez-Hernández and Seira, 2021), which means that the assumptions made in Section 6 apply. Depending on the time frame in consideration and units of the good, consumers might either have unit demand, a common assumption in the empirical discrete choice models of papers cited in Section 1, or continuous demand, which would require extending the analysis of this paper.

In closing, it is important to point out that policy for the public option in any real-world market must ultimately be informed by careful empirical research. While the focus of this paper has been a theoretical, rather than empirical, analysis, many of the theoretical results were derived with an eye toward future integration with empirical estimates. In particular, Proposition 4 and Theorem 2 show how the optimal mechanism can be explicitly computed given the distribution of consumer preferences, while Proposition 2 suggests an algorithm by which the value of the public option can be evaluated quickly.

8 Conclusion

The outcome, and hence success, of any public program depends not only on its direct effects, but also on its indirect effects that obtain only in equilibrium. Almost tautologically, any analysis of equilibrium effects requires the designer to solve a fixed-point problem. A key contribution of this paper is to provide a tractable mechanism design framework for this analysis.

This framework produces a simple characterization of the optimal way for a designer to provide a public option when consumers can also purchase the good in a private market. It also yields an explicit algorithm to compute the optimal mechanism. The optimality of rationing depends on the amount of inequality among consumers, the cost of providing the public option relative to the price in the private market, and how much of the public option is sold relative to the private good. The optimal mechanism simplifies when the designer can tax the private good or impose an individual mandate, but in general rationing remains optimal.

The simplicity of optimal mechanisms identified in this paper relies on a few key assumptions: the linearity of consumer utility subject to unit demand, the absence of selection effects, and the competitiveness of the private market. Even if these assumptions are relaxed, however, the economic intuitions of the present framework are likely to extend to richer settings. Of these, the

easiest assumption to relax is the linearity of consumer utility: using the generalized concavification approach (cf. Proposition 4 and Theorem 2) of Section 4, the insights of this paper extend to a model in which consumers have quadratic utility.

It is important to also emphasize the limitations of the present analysis. For one, this paper abstracts away from political economy considerations, even though redistributive policies—optimal or not—often require political support; one must look no further than the "Medicare for All" debate for evidence. For another, externalities and behavioral reasons have often been cited as possible justification for the public provision of private goods. The extension of the framework presented in this paper to these richer environments is left for future research.

References

AKBARPOUR, M. ® P. DWORCZAK ® S. D. KOMINERS (2020): "Redistributive Allocation Mechanisms," *Working paper*.

ATAL, J. P., J. I. CUESTA, F. GONZÁLEZ, AND C. OTERO (2021): "The Economics of the Public Option: Evidence from Local Pharmaceutical Markets," *Working paper*.

ATKINSON, A. B. AND J. E. STIGLITZ (1976): "The Design of Tax Structure: Direct Versus Indirect Taxation," *Journal of Public Economics*, 6, 55–75.

AUMANN, R. J. AND M. MASCHLER (1995): *Repeated Games with Incomplete Information*, Cambridge, MA: MIT Press.

BAUER, H. (1958): "Minimalstellen von Funktionen und Extremalpunkte," *Archiv der Mathematik*, 9, 389–393.

BESLEY, T. AND S. COATE (1991): "Public Provision of Private Goods and the Redistribution of Income," *American Economic Review*, 81, 979–984.

BLACKORBY, C. AND D. DONALDSON (1988): "Cash Versus Kind, Self-Selection, and Efficient Transfers," *American Economic Review*, 691–700.

BONNANS, J. F. AND A. SHAPIRO (2000): *Perturbation Analysis of Optimization Problems*, New York, NY: Springer.

BOYD, S. AND L. VANDENBERGHE (2004): *Convex Optimization*, Cambridge, UK: Cambridge University Press.

BULOW, J. AND J. ROBERTS (1989): "The Simple Economics of Optimal Auctions," *Journal of Political Economy*, 97, 1060–1090.

CARROLL, G. AND I. SEGAL (2019): "Robustly Optimal Auctions with Unknown Resale Opportunities," *Review of Economic Studies*, 86, 1527–1555.

CHE, Y.-K., W. DESSEIN, AND N. KARTIK (2013a): "Pandering to Persuade," *American Economic Review*, 103, 47–79.

CHE, Y.-K., I. GALE, AND J. KIM (2013b): "Assigning Resources to Budget-Constrained Agents," *Review of Economic Studies*, 80, 73–107.

CONDORELLI, D. (2013): "Market and Non-Market Mechanisms for the Optimal Allocation of Scarce Resources," *Games and Economic Behavior*, 82, 582–591.

CURRIE, J. AND F. GAHVARI (2008): "Transfers in Cash and In-Kind: Theory Meets the Data," *Journal of Economic Literature*, 46, 333–83.

DIAMOND, P. A. AND J. A. MIRRLEES (1971): "Optimal Taxation and Public Production I: Production Efficiency," *American Economic Review*, 61, 8–27.

DIAMOND, R., T. McQUADE, AND F. QIAN (2019): "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco," *American Economic Review*, 109, 3365–94.

VAN DIJK, W. (2019): "The Socio-Economic Consequences of Housing Assistance," *Working paper*.

DINERSTEIN, M., C. NEILSON, AND S. OTERO (2020): "The Equilibrium Effects of Public Provision in Education Markets: Evidence from a Public School Expansion Policy," *Working paper*.

DOVAL, L. AND V. SKRETA (2021): "Constrained Information Design," *Working paper*.

DWORCZAK, P. (2020): "Mechanism Design With Aftermarkets: Cutoff Mechanisms," *Econometrica*, 88, 2629–2661.

DWORCZAK, P. ® S. D. KOMINERS ® M. AKBARPOUR (2021): "Redistribution through Markets," *Econometrica*, 89, 1665–1698.

EINAV, L. AND A. FINKELSTEIN (2011): "Selection in Insurance Markets: Theory and Empirics in Pictures," *Journal of Economic Perspectives*, 25, 115–38.

EPPLE, D. AND R. E. ROMANO (1996): "Public Provision of Private Goods," *Journal of Political Economy*, 104, 57–84.

FUCHS, W. AND A. SKRZYPACZ (2015): "Government Interventions in a Dynamic Market with Adverse Selection," *Journal of Economic Theory*, 158, 371–406.

GAHVARI, F. AND E. MATTOS (2007): "Conditional Cash Transfers, Public Provision of Private Goods, and Income Redistribution," *American Economic Review*, 97, 491–502.

GERUSO, M. AND T. J. LAYTON (2017): "Selection in Health Insurance Markets and Its Policy Remedies," *Journal of Economic Perspectives*, 31, 23–50.

HANCHE-OLSEN, H. AND H. HOLDEN (2010): "The Kolmogorov–Riesz Compactness Theorem," *Expositiones Mathematicae*, 28, 385–394.

HARTLINE, J. D. (2013): "Mechanism Design and Approximation," *Book draft*.

HOTELLING, H. (1931): "The Economics of Exhaustible Resources," *Journal of Political Economy*, 39, 137–175.

JEHIEL, P. AND B. MOLDOVANU (2001): "Efficient Design with Interdependent Valuations," *Econometrica*, 69, 1237–1259.

JIMÉNEZ-HERNÁNDEZ, D. AND E. SEIRA (2021): "Should the Government Sell You Goods? Evidence from the Milk Market in Mexico," *Working paper*.

KAMENICA, E. AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

KLEINER, A., B. MOLDOVANU, AND P. STRACK (2021): "Extreme Points and Majorization: Economic Applications," *Econometrica*, 89, 1557–1593.

LE TREUST, M. AND T. TOMALA (2019): "Persuasion with Limited Communication Capacity," *Journal of Economic Theory*, 184, 104940.

LOERTSCHER, S. AND E. V. MUIR (2020): "Monopoly Pricing, Optimal Rationing, and Resale," *Working paper*.

MANELLI, A. M. AND D. R. VINCENT (2007): "Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly," *Journal of Economic Theory*, 137, 153–185.

MASKIN, E. AND J. RILEY (1984): "Monopoly with Incomplete Information," *RAND Journal of Economics*, 15, 171–196.

MILGROM, P. AND I. SEGAL (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70, 583–601.

MUNRO, A. (1991): "The Optimal Public Provision of Private Goods," *Journal of Public Economics*, 44, 239–261.

Mussa, M. and S. Rosen (1978): "Monopoly and Product Quality," *Journal of Economic Theory*, 18, 301–317.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

Philippon, T. and V. Skreta (2012): "Optimal Interventions in Markets with Adverse Selection," *American Economic Review*, 102, 1–28.

Royden, H. L. and P. M. Fitzpatrick (2010): *Real Analysis*, Boston, MA: Prentice Hall, 4th edition.

Saez, E. and S. Stantcheva (2016): "Generalized Social Marginal Welfare Weights for Optimal Tax Theory," *American Economic Review*, 106, 24–45.

Samuelson, W. (1984): "Bargaining under Asymmetric Information," *Econometrica*, 54, 995–1005.

Sitaraman, G. and A. L. Alstott (2019): *The Public Option: How to Expand Freedom, Increase Opportunity, and Promote Equality*, Cambridge, MA: Harvard University Press.

Skreta, V. (2006): "Sequentially Optimal Mechanisms," *Review of Economic Studies*, 73, 1085–1111.

Summers, L. H. (1989): "Some Simple Economics of Mandated Benefits," *American Economic Review: Papers & Proceedings*, 79, 177–183.

Szapiel, W. (1975): "Points Extrémaux dans les Ensembles Convexes (I). Théorie Générale," *Bulletin de l'Académie Polonaise des Sciences. Série des Sciences Mathématiques, Astronomiques et Physiques*, 23, 939–945.

Tirole, J. (2012): "Overcoming Adverse Selection: How Public Intervention Can Restore Market Functioning," *American Economic Review*, 102, 29–59.

Toikka, J. (2011): "Ironing without Control," *Journal of Economic Theory*, 146, 2510–2526.

Weitzman, M. L. (1977): "Is the Price System or Rationing More Effective in Getting a Commodity to Those Who Need It Most?" *Bell Journal of Economics*, 8, 517–524.

Appendix A Omitted proofs

A.1 Proofs of results in Section 3

A.1.1 Proof of Lemma 1

For any given price p of the private good, (IC) can be written as

$$\left[\delta\theta/v - (\theta/v - p)_+\right] X(\delta, \theta, v) - T(\delta, \theta, v)$$
$$\geq \left[\delta\theta/v - (\theta/v - p)_+\right] X(\delta', \theta', v') - T(\delta', \theta', v') \quad \text{for any } (\delta, \theta, v) \text{ and } (\delta', \theta', v').$$

This implies that, for any (δ, θ, v) and (δ', θ', v'):

$$\left\{\left[\delta\theta/v - (\theta/v - p)_+\right] - \left[\delta'\theta'/v' - (\theta'/v' - p)_+\right]\right\} \cdot \left[X(\delta, \theta, v) - X(\delta', \theta', v')\right] \geq 0.$$

Hence, for any (δ, θ, v) and (δ', θ', v'):

$$\left[\delta\theta/v - (\theta/v - p)_+\right] > \left[\delta'\theta'/v' - (\theta'/v' - p)_+\right] \implies X(\delta, \theta, v) \geq X(\delta', \theta', v').$$

Define $y(\eta, \theta, v) \equiv X\left(v/\theta \cdot \left(\eta + (\theta/v - p)_+\right), \theta, v\right)$. For sufficiently small $\varepsilon > 0$, the above implies that

$$y(\eta + \varepsilon, \theta, v) \geq y(\eta, \theta', v') \quad \text{for any } \theta, \theta', v, v' \text{ and almost every } \eta \in [\underline{\eta}, \overline{\eta}].$$

This shows that $y(\cdot, \theta, v)$ is increasing for each θ, v over the interval $[\underline{\eta}, \overline{\eta}]$; hence it is continuous almost everywhere. Taking $\varepsilon \to 0$ yields

$$y(\eta, \theta, v) \geq y(\eta, \theta', v') \quad \text{for any } \theta, \theta', v, v' \text{ and almost every } \eta \in [\underline{\eta}, \overline{\eta}].$$

Therefore, $X(\delta, \theta, v) = x(\delta r - (r - p^*)_+)$, where x is an increasing function, for almost every $\delta r - (r - p^*)_+ \in [\underline{\eta}, \overline{\eta}]$. This proves statement *(i)* of the lemma; statement *(ii)* follows directly from the envelope theorem of Milgrom and Segal (2002).

A.1.2 Proof of Lemma 2

Suppose that the allocation function x effects the price p in the private market. Because the private market is competitive, $D(p) = S(p)$. Conversely, observe that $D(\cdot)$ is decreasing and continuous (by continuity of the integral), and $S(\cdot)$ is strictly increasing and continuous (by assumption).

Moreover, $S(\underline{r}) = 0 \le D(\underline{r})$ and $S(\overline{r}) > S(\underline{r}) = 0 = D(\overline{r})$. Hence there is a unique price p in the interval $[\underline{r}, \overline{r}]$ at which $D(p) = S(p)$.

A.1.3 Proof of Lemma 3

Suppose that $x \stackrel{\text{a.e.}}{=} \widetilde{x}$ for some increasing function $\widetilde{x} : [\underline{\eta}, \overline{\eta}] \to [0, 1]$ satisfying $\operatorname{im} \widetilde{x} \subset \{0, 1\}$, and that $x = \alpha x_1 + (1 - \alpha) x_2$ for $x_1, x_2 \in K$ and $\alpha \in (0, 1)$. Since $x_1, x_2 \in K$, there exist increasing functions $\widetilde{x}_1, \widetilde{x}_2 : [\underline{\eta}, \overline{\eta}] \to [0, 1]$ such that $x_1 \stackrel{\text{a.e.}}{=} \widetilde{x}_1$ and $x_2 \stackrel{\text{a.e.}}{=} \widetilde{x}_2$. Then $\alpha \widetilde{x}_1(\eta) + (1 - \alpha) \widetilde{x}_2(\eta) = \widetilde{x}(\eta) \in \{0, 1\}$ for almost every $\eta \in [\underline{\eta}, \overline{\eta}]$, which implies that $\widetilde{x}_1(\eta) = \widetilde{x}_2(\eta) = \widetilde{x}(\eta)$ for almost every $\eta \in [\underline{\eta}, \overline{\eta}]$. Thus $x_1 = x_2$; hence $x \in \operatorname{ex} K$.

Conversely, let $x \stackrel{\text{a.e.}}{=} \widetilde{x}$ for some increasing function $\widetilde{x} : [\underline{\eta}, \overline{\eta}] \to [0, 1]$, where the set $\Gamma = \{\eta \in [\underline{\eta}, \overline{\eta}] : \widetilde{x}(\eta) \notin \{0, 1\}\}$ has positive measure. Define $\widetilde{x}_1, \widetilde{x}_2 : [\underline{\eta}, \overline{\eta}] \to [0, 1]$ by $\widetilde{x}_1 = \widetilde{x}^2$ and $\hat{x}_2 = 2\widetilde{x} - \widetilde{x}^2$; by construction, $\widetilde{x}_1$ and $\widetilde{x}_2$ are increasing and $\widetilde{x} = (\widetilde{x}_1 + \widetilde{x}_2) / 2$. Let $x_1, x_2 \in K$ be the equivalence classes of functions satisfying $x_1 \stackrel{\text{a.e.}}{=} \widetilde{x}_1$ and $x_2 \stackrel{\text{a.e.}}{=} \widetilde{x}_2$; since $\widetilde{x}_1(\eta) \ne \widetilde{x}_2(\eta)$ on Γ (which has positive measure by assumption) it follows that $x_1 \ne x_2$. Therefore $x = (x_1 + x_2) / 2$ where $x_1, x_2 \in K$ are distinct; hence $x \notin \operatorname{ex} K$.

A.1.4 Completion of the proof of Theorem 1

To apply the constrained maximum principle, it remains to show that K is convex and compact. Moreover, the proof of Theorem 1 assumes the existence of an optimal mechanism; and it remains to show that an optimal mechanism exists.

Convexity of K. Let $x_1, x_2 \in K$ with $x_1 \stackrel{\text{a.e.}}{=} \widetilde{x}_1$ and $x_2 \stackrel{\text{a.e.}}{=} \widetilde{x}_2$ for increasing functions $\widetilde{x}_1, \widetilde{x}_2 : [\underline{\eta}, \overline{\eta}] \to [0, 1]$. For any $\alpha \in [0, 1]$, observe that $\alpha x_1 + (1 - \alpha) x_2 \stackrel{\text{a.e.}}{=} \alpha \widetilde{x}_1 + (1 - \alpha) \widetilde{x}_2$ and that $\alpha \widetilde{x}_1 + (1 - \alpha) \widetilde{x}_2 : [\underline{\eta}, \overline{\eta}] \to [0, 1]$ is increasing. Thus K is convex.

Compactness of K. Observe that $L^1([\underline{\eta}, \overline{\eta}])$ is complete by the Riesz–Fisher theorem (cf. Section 7.3 of Royden and Fitzpatrick, 2010); thus it suffices to show that K is closed and totally bounded. By the Kolmogorov–Riesz theorem (cf. Theorem 5 of Hanche-Olsen and Holden, 2010), K is totally bounded. Let $x_n \stackrel{L^1}{\to} x$; passing to a subsequence if necessary, $x_n \to x$ pointwise. By the bounded convergence theorem, it follows that K is closed. Thus K is compact.

Existence of an optimal mechanism. Say that an effective mechanism is *feasible* if it satisfies (IC), (IR), the market-clearing constraint (MC), and the budget constraint. Let $\Pi \subset [\underline{r}, \overline{r}]$ be the

set of prices in the private market that can be effected by a feasible effective mechanism. Observe that the competitive price p_{c} obtains in the private market if the designer does not intervene (which is always feasible), so Π is nonempty.

For any sequence $\{p_n\} \subset \Pi$ converging to p, let $\{(x_n, t_n)\}$ be a sequence of mechanisms that are optimal for the price p_n (note that an optimal mechanism conditional on any price $p_n \in \Pi$ always exists). Let (x, t) be the mechanism obtained by taking the pointwise limit of $\{(x_n, t_n)\}$; such a limit exists (by passing to a subsequence if necessary) because $0 \le x_n \le 1$ and t_n is bounded as a result of the envelope theorem and the budget constraint. It is straightforward to verify that (x, t) is feasible. Because $S(\cdot)$ is continuous, (x, t) effects the price p in the private market: $p \in \Pi$.

Let $\Omega^*(\pi)$ denote the optimal value of the design problem conditional on the price $\pi \in \Pi$ in the private market. By continuity of the design objective, the above argument shows that $\limsup_{p_n \to p} \Omega^*(p_n) \le \Omega^*(p)$. Therefore, the supremum of $\Omega^*(\cdot)$ must be attained on Π; hence an optimal mechanism exists.

A.2 Proofs of results in Section 4

A.2.1 Proof of Proposition 2

The result follows from strong duality, which holds in this setting because K is nonempty, compact, and convex by the proof of Theorem 1 (cf. Theorem 2.202 of Bonnans and Shapiro, 2000). It is well-known that the dual program is convex (cf. Section 5.2 of Boyd and Vandenberghe, 2004).

A.2.2 Proof of Proposition 3

(i) Let $c_1 > c_2$ and suppose that the optimal mechanism is (x_1, t_1) when the marginal cost of the public option is c_1. Then (x_1, t_1) must still satisfy the budget constraint and effect the same price in the private market when the marginal cost is c_2. Since neither the design objective nor h_{c} depends on the marginal cost, the value of the public option when the marginal cost is c_2 must be (weakly) higher.

(ii) Suppose that there is uniformly more wealth associated with $\lambda_1(r, \delta)$ than $\lambda_2(r, \delta)$, and let (x_1, t_1) be the optimal mechanism for the consumer population with Pareto weights $\lambda_1(r, \delta)$. Suppose that (x_1, t_1) effects the price p_1 in the private market. Because the constraints of the design problem do not depend on the Pareto weights, (x_1, t_1) remains a feasible mechanism at the price of p_1 when consumers have Pareto weights $\lambda_2(r, \delta)$. Then the change in the

value of the public option when consumers have Pareto weights $\lambda_2(r, \delta)$ is

$$(\underline{\Lambda}_2 - \underline{\Lambda}_1) \cdot (\underline{U} - B) + \int_{\underline{\eta}}^{\overline{\eta}} [\Lambda_2(\eta) - \Lambda_1(\eta)] \, x_1(\eta) \, \mathrm{d}\eta$$

$$+ \int_{\underline{r}}^{\overline{r}} \int_{\underline{\delta}}^{\overline{\delta}} \left[(r - p_1)_+ - (r - p_\mathrm{c})_+ \right] [\lambda_2(r, \delta) - \lambda_1(r, \delta)] \ \mathrm{d}F(r, \delta).$$

By assumption, $\lambda_1(r, \delta) \leq \lambda_2(r, \delta)$ for every (r, δ). Since $p_1 \leq p_\mathrm{c}$, this expression must be non-negative for sufficiently small values of B.

(iii) Let $\lambda_1(r, \delta)$ be more unequal than $\lambda_2(r, \delta)$, and denote the corresponding effective survival Pareto functions by $\Lambda_1(\eta)$ and $\Lambda_2(\eta)$. If η is decreasing in r, then

$$\Lambda(t) = \int_t^{\overline{\eta}} \lambda(s) \, \mathrm{d}G(s) = \mathbf{E}_r \left\{ \mathbf{E}_\delta \left[\lambda(r, \delta) \,|\, \delta r - (r - p)_+ \geq t \right] \right\} \cdot \mathbf{P} \left[\delta r - (r - p)_+ \geq t \right]$$

$$= \mathbf{E}_r \left\{ \mathbf{E}_\delta \left[\lambda(r, \delta) \,|\, r \leq \eta^{-1}(t, \delta) \right] \right\} \cdot \mathbf{P} \left[\delta r - (r - p)_+ \geq t \right].$$

Thus $\Lambda_1(\eta) \geq \Lambda_2(\eta)$ for every η. Let (x_1, t_1) be the optimal mechanism for the consumer population with Pareto weights $\lambda_1(r, \delta)$, and suppose that (x_1, t_1) effects the price p_1 in the private market. As in the proof of (ii) above, the change in the value of the public option when consumers have Pareto weights $\lambda_2(r, \delta)$ is

$$\underbrace{\int_{\underline{\eta}}^{\overline{\eta}} [\Lambda_2(\eta) - \Lambda_1(\eta)] \, x_1(\eta) \, \mathrm{d}\eta}_{\leq 0} + \int_{\underline{r}}^{\overline{r}} \int_{\underline{\delta}}^{\overline{\delta}} \left[(r - p_1)_+ - (r - p_\mathrm{c})_+ \right] [\lambda_2(r, \delta) - \lambda_1(r, \delta)] \ \mathrm{d}F(r, \delta).$$

For any p, observe that

$$\int_{\underline{r}}^{\overline{r}} \int_{\underline{\delta}}^{\overline{\delta}} (r - p)_+ \, [\lambda_2(r, \delta) - \lambda_1(r, \delta)] \ \mathrm{d}F(r, \delta) = \int_p^{\overline{r}} \mathbf{E}_{(r, \delta)} \left[[\lambda_1(r, \delta) - \lambda_2(r, \delta)] \cdot \mathbf{1}_{r \geq s} \right] \ \mathrm{d}s.$$

Thus

$$\int_{\underline{r}}^{\overline{r}} \int_{\underline{\delta}}^{\overline{\delta}} \left[(r - p_1)_+ - (r - p_\mathrm{c})_+ \right] [\lambda_2(r, \delta) - \lambda_1(r, \delta)] \ \mathrm{d}F(r, \delta)$$

$$= \int_{p_1}^{p_\mathrm{c}} \mathbf{E}_{(r, \delta)} \left[[\lambda_1(r, \delta) - \lambda_2(r, \delta)] \cdot \mathbf{1}_{r \geq s} \right] \ \mathrm{d}s \leq 0.$$

A.3 Proofs of results in Section 5

A.3.1 Proof of Proposition 4

The proof of Proposition 4 makes use of the following lemma:

Lemma 4. Fix $\varepsilon > 0$ and the Lagrange multipliers $\mu_\varepsilon \in \mathbb{R}$ and $\alpha_\varepsilon \geq 0$. Then there exists a unique solution

$$x_\varepsilon^\circ \in \arg\max_{x \in K} \int_{\underline{\eta}}^{\overline{\eta}} \left\{ \left[\frac{\Lambda(\eta)}{g(\eta)} + \mu_\varepsilon \phi(\eta) + \alpha_\varepsilon \left[\eta - c - \frac{1 - G(\eta)}{g(\eta)} \right] \right] x(\eta) - \varepsilon \left[x(\eta) \right]^2 \right\} \, dG(\eta).$$

Proof of Lemma 4. View each effective allocation function x as an element of the Hilbert space $L^2([0,1]; G)$, and recall from the proof of Theorem 1 that $K \subset L^2([0,1]; G)$ is nonempty, compact, and convex. Moreover, by completing the square,

$$x_\varepsilon^\circ \in \arg\min_{x \in K} \int_{\underline{\eta}}^{\overline{\eta}} \left\{ x(\eta) - \frac{1}{\varepsilon} \left[\frac{\Lambda(\eta)}{g(\eta)} + \mu_\varepsilon \phi(\eta) + \alpha_\varepsilon \left[\eta - c - \frac{1 - G(\eta)}{g(\eta)} \right] \right] \right\}^2 \, dG(\eta).$$

Thus x_ε° is the unique projection of

$$\frac{1}{\varepsilon} \left[\frac{\Lambda(\eta)}{g(\eta)} + \mu_\varepsilon \phi(\eta) + \alpha_\varepsilon \left[\eta - c - \frac{1 - G(\eta)}{g(\eta)} \right] \right] \quad \text{onto } K. \qquad \square$$

To prove Proposition 4, begin by assuming that $0 < x_\varepsilon(t) < 1$ for all $t \in [0,1]$. I claim that x_ε is in fact the projection of

$$\psi_\varepsilon(\eta) \equiv \frac{1}{\varepsilon} \left[\frac{\Lambda(\eta)}{g(\eta)} + \mu_\varepsilon^* \phi(\eta) + \alpha_\varepsilon^* \left[\eta - c - \frac{1 - G(\eta)}{g(\eta)} \right] \right]$$

onto the set of all increasing functions $I \equiv \{x : [0,1] \to \mathbb{R} \text{ such that } x \text{ is increasing}\}$. By Lemma 4, this projection exists and is unique; denote this projection by x_ε^*. By the definition of the projection operator Π_I,

$$\int_{\underline{\eta}}^{\overline{\eta}} \left[x_\varepsilon^*(\eta) - \psi_\varepsilon(\eta) \right] v(\eta) \, dG(\eta) \geq 0 \quad \text{for any } v \in I.$$

In particular, choose $v(s) = \mathbf{1}_{s > t}$ for some $t \in [0,1]$. Then

$$L_\varepsilon(\eta) \equiv \int_\eta^{\overline{\eta}} x_\varepsilon^*(s) \, dG(s) \geq \int_\eta^{\overline{\eta}} \psi_\varepsilon(s) \, dG(s) = H_\varepsilon(\eta).$$

40

Since x_ε^* is increasing on $[0, 1]$, L_ε must be concave; hence L_ε is a concave function that bounds H_ε from above. Now, if L_ε is not the pointwise smallest such function, then there exists x_ε^{**} that coincides with x_ε^* everywhere except on some interval $(\eta_1, \eta_2) \subset [\underline{\eta}, \overline{\eta}]$ on which x_ε^{**} is constant. Then integration by parts yields the following contradiction:

$$0 \le \int_{\underline{\eta}}^{\overline{\eta}} [x_\varepsilon^*(\eta) - \psi_\varepsilon(\eta)] [x_\varepsilon^{**}(\eta) - x_e^*(\eta)] \, \mathrm{d}G(\eta)$$

$$= -\int_{\underline{\eta}}^{\overline{\eta}} [H_\varepsilon(\eta) - L_\varepsilon(\eta)] \, \mathrm{d} [x_\varepsilon^{**}(\eta) - x_\varepsilon^*(\eta)] = \int_{\eta_1}^{\eta_2} [H_\varepsilon(\eta) - L_\varepsilon(\eta)] \, \mathrm{d}x_\varepsilon^*(\eta) < 0.$$

Therefore $L_\varepsilon = \mathrm{co}\, H_\varepsilon$ must be the concave closure of H_ε.

It remains to consider the case where either $x_\varepsilon(\eta) \le 0$ or $1 \le x_\varepsilon(\eta)$ (or both) for some $\eta \in [\underline{\eta}, \overline{\eta}]$. Define η_1, η_2 so that $x_\varepsilon(\eta) \le 0$ for all $\eta \in [\underline{\eta}, \eta_1)$; $0 < x_\varepsilon(\eta) < 1$ for all $\eta \in (\eta_1, \eta_2)$; and $1 \le x_\varepsilon(\eta)$ for all $\eta \in (\eta_2, \overline{\eta}]$. The argument above shows that

$$x_\varepsilon \in \arg\min_{x \in K} \int_{\eta_1}^{\eta_2} [x(\eta) - \psi_\varepsilon(\eta)] \, \mathrm{d}G(\eta).$$

Because $K \subset I$, the argument above also shows that the constraint $x(\eta) \ge 0$ must bind for $t \in [\underline{\eta}, \eta_1]$; and the constraint $x(\eta) \le 1$ must bind for $\eta \in [\eta_2, \overline{\eta}]$. This yields the solution x_ε^* in the statement of Proposition 4.

A.3.2 Proof of Theorem 2

Since $\{x_\varepsilon^*\}$ is uniformly bounded, Helly's selection theorem applies: there exists a subsequence of $\{x_\varepsilon^*\}$ that converges pointwise to some function x^*. Since convergence is pointwise and each x_ε^* in the subsequence is feasible, hence x^* must also be feasible.

Suppose on the contrary that x^* is not a solution of the design problem. Then there exists an effective allocation function x° that satisfies (MC) and the budget constraint, such that

$$\int_{\underline{\eta}}^{\overline{\eta}} \frac{\Lambda(\eta)}{g(\eta)} \cdot x^\circ(\eta) \, \mathrm{d}G(\eta) > \int_{\underline{\eta}}^{\overline{\eta}} \frac{\Lambda(\eta)}{g(\eta)} \cdot x^*(\eta) \, \mathrm{d}G(\eta).$$

However, continuity implies the existence of some $\varepsilon > 0$ such that

$$\int_{\underline{\eta}}^{\overline{\eta}} \left[\frac{\Lambda(\eta)}{g(\eta)} \cdot x^\circ(\eta) - \frac{\varepsilon}{2} [x^\circ(\eta)]^2 \right] \, \mathrm{d}G(\eta) > \int_{\underline{\eta}}^{\overline{\eta}} \left[\frac{\Lambda(\eta)}{g(\eta)} \cdot x_\varepsilon^*(\eta) - \frac{\varepsilon}{2} [x_\varepsilon^*(\eta)]^2 \right] \, \mathrm{d}G(\eta).$$

This contradicts the optimality of x_ε^* for the respective augmented problem.

It remains to show that the entire sequence of $\{x_\varepsilon^*\}$ (*i.e.*, not just a subsequence) converges to x^*. This follows from the continuity of x_ε^* in ε due to the continuity of the projection operator.

A.4   Proofs of results in Section 6

A.4.1   Proof of Proposition 5

I abuse notation by writing the mechanism as (x, t), where the allocation function $x : [p, \bar{r}] \to [0, 1]$ and the transfer function $t : [p, \bar{r}] \to \mathbb{R}$ are functions of each consumer's rate of substitution r between the private good and money. By Lemma 1, $x(r)$ is decreasing on $[p, \bar{r}]$ and can be extended to the interval $[\underline{r}, p]$ via

$$x(r) = x\left(\frac{p - \delta r}{1 - \delta}\right) \quad \text{for } r \in [\underline{r}, p].$$

Similarly, $t(r)$ can be extended to the interval $[\underline{r}, p]$ via

$$t(r) = t\left(\frac{p - \delta r}{1 - \delta}\right) \quad \text{for } r \in [\underline{r}, p].$$

The envelope theorem yields

$$[p - (1 - \delta) r] x(r) - t(r) = \underline{U} + (1 - \delta) \int_r^{\bar{r}} x(s) \, \mathrm{d}s \quad \text{for } r \in [p, \bar{r}].$$

Thus weighted consumer surplus for consumers with $r \in [p, \bar{r}]$ is

$$\underline{U} \cdot \Lambda(p) + (1 - \delta) \int_p^{\bar{r}} [\Lambda(p) - \Lambda(r)] x(r) \, \mathrm{d}r + \int_p^{\bar{r}} (r - p) \, \mathrm{d}F(r).$$

Similarly, weighted consumer surplus for consumers with $r \in [\underline{r}, p]$ is

$$\underline{U} \cdot [\underline{\Lambda} - \Lambda(p)] + (1 - \delta) \left\{ \int_p^{\frac{p - \delta r}{1 - \delta}} \left[\Lambda\left(\frac{p - (1 - \delta) r}{\delta}\right) - \Lambda(p) \right] x(r) \, \mathrm{d}r + \int_{\frac{p - \delta r}{1 - \delta}}^{\bar{r}} [\underline{\Lambda} - \Lambda(p)] x(r) \, \mathrm{d}r \right\}.$$

Then weighted total consumer surplus is

$$\underline{U} \cdot \underline{\Lambda} + \int_p^{\bar{r}} (r - p) \, \mathrm{d}F(r) + (1 - \delta) \int_p^{\bar{r}} \hat{\Lambda}(r) x(r) \, \mathrm{d}r.$$

Here,

$$
\hat{\Lambda}(r) \equiv
\begin{cases}
\Lambda\left(\dfrac{p - (1 - \delta)\, r}{\delta}\right) - \Lambda(r) & \text{for } p \le r \le \dfrac{p - \delta \underline{r}}{1 - \delta}, \\[3mm]
\underline{\Lambda} - \Lambda(r) & \text{for } \dfrac{p - \delta \underline{r}}{1 - \delta} \le r \le \bar{r}.
\end{cases}
$$

Now consider the profit from supplying the public option. The profit from consumers with $r \in [p, \bar{r}]$ is

$$
-\underline{U} \cdot [1 - F(p)] + \int_p^{\bar{r}} \left\{ p - c - (1 - \delta)\left[r + \frac{F(r) - F(p)}{f(r)} \right] \right\} x(r) \, \mathrm{d}F(r).
$$

Similarly, the profit from consumers with $r \in [\underline{r}, p]$ is

$$
-\underline{U} \cdot F(p) + \int_p^{\frac{p - \delta \underline{r}}{1 - \delta}} [p - c - (1 - \delta)\, r]\, x(r) \cdot \frac{1 - \delta}{\delta} f\left(\frac{p - (1 - \delta)\, r}{\delta} \right) \, \mathrm{d}r
$$

$$
- \int_p^{\frac{p - \delta \underline{r}}{1 - \delta}} (1 - \delta) \left[F(p) - F\left(\frac{p - (1 - \delta)\, r}{\delta} \right) \right] x(r) \, \mathrm{d}r - \int_{\frac{p - \delta \underline{r}}{1 - \delta}}^{\bar{r}} (1 - \delta)\, F(p) x(r) \, \mathrm{d}r.
$$

Then the designer's total profit from the public option is

$$
-\underline{U} + \int_p^{\bar{r}} [p - c - (1 - \delta)\, r]\, x(r) \hat{f}(r) \, \mathrm{d}r - (1 - \delta) \int_p^{\bar{r}} \hat{F}(r) x(r) \, \mathrm{d}r.
$$

Here,

$$
\hat{f}(r) \equiv
\begin{cases}
f(r) + \dfrac{1 - \delta}{\delta} f\left(\dfrac{p - (1 - \delta)\, r}{\delta} \right) & \text{for } p \le r \le \dfrac{p - \delta \underline{r}}{1 - \delta}, \\[3mm]
f(r) & \text{for } \dfrac{p - \delta \underline{r}}{1 - \delta} < r \le \bar{r},
\end{cases}
$$

$$
\hat{F}(r) \equiv
\begin{cases}
F(r) - F\left(\dfrac{p - (1 - \delta)\, r}{\delta} \right) & \text{for } p \le r \le \dfrac{p - \delta \underline{r}}{1 - \delta}, \\[3mm]
F(r) & \text{for } \dfrac{p - \delta \underline{r}}{1 - \delta} \le r \le \bar{r}.
\end{cases}
$$

Thus the design problem is

$$
\max_{x,\overline{U}} \ V(p) + \underline{U} \cdot \underline{\Lambda} + \int_p^{\overline{r}} (r - p) \ \mathrm{d}F(r) + (1 - \delta) \int_p^{\overline{r}} \hat{\Lambda}(r) x(r) \ \mathrm{d}r
$$

$$
\text{s.t.} \begin{cases}
x : [p, \overline{r}] \to [0, 1] \text{ is decreasing;} \ \underline{U} \geq 0, \\[2mm]
\displaystyle\int_p^{\overline{r}} [1 - x(r)] \ \mathrm{d}F(r) = S(p), \\[3mm]
\displaystyle B - \underline{U} + \int_p^{\overline{r}} [p - c - (1 - \delta) r] \, x(r) \hat{f}(r) \ \mathrm{d}r - (1 - \delta) \int_p^{\overline{r}} \hat{F}(r) x(r) \ \mathrm{d}r \geq 0.
\end{cases}
$$

Substitute $y(r) = 1 - x(r)$ and rewrite the design problem as follows in terms of y:

$$
\max_{y \text{ increasing}} \ -\int_p^{\overline{r}} \hat{\Lambda}(r) y(r) \ \mathrm{d}r
$$

$$
\text{s.t.} \begin{cases}
\displaystyle\int_p^{\overline{r}} y(r) \ \mathrm{d}F(r) = S(p), \\[3mm]
\displaystyle\int_p^{\overline{r}} \left\{ p - c - (1 - \delta) \left[r + \frac{\hat{F}(r)}{\hat{f}(r)} \right] \right\} y(r) \cdot \frac{\hat{f}(r)}{f(r)} \ \mathrm{d}F(r) \leq B_0.
\end{cases}
$$

Suppose that, instead of a budget constraint, $\overline{U}$ is constrained to be in the interval $[0, \kappa]$. Then the design problem can be written as

$$
\max_{y \text{ increasing}} \left\{ -\int_p^{\overline{r}} \hat{\Lambda}(r) y(r) \ \mathrm{d}r : \int_p^{\overline{r}} y(r) \ \mathrm{d}F(r) = S(p) \right\}.
$$

Because $f(r)$ is decreasing, observe that

$$
\frac{\mathrm{d}}{\mathrm{d}r} \left[-\frac{\hat{\Lambda}(r)}{f(r)} \right] = -\frac{\hat{\Lambda}'(r)}{f(r)} + \underbrace{\frac{\hat{\Lambda}(r) f'(r)}{[f(r)]^2}}_{\leq 0} \leq \begin{cases}
\dfrac{\Lambda'(r) + \frac{1-\delta}{\delta} \Lambda' \left(\frac{p - (1-\delta)r}{\delta} \right)}{f(r)} & \text{for } p \leq r \leq \dfrac{p - \delta r}{1 - \delta}, \\[5mm]
\dfrac{\Lambda'(r)}{f(r)} & \text{for } \dfrac{p - \delta r}{1 - \delta} < r \leq \overline{r}.
\end{cases}
$$

Because $\Lambda'(r) \leq 0$ for all r and $\hat{\Lambda}(r)$ is continuous at $r = (p - \delta r) / (1 - \delta)$, thus $-\hat{\Lambda}(r)/f(r)$ is decreasing in r on $[p, \overline{r}]$. Hence the optimal mechanism sets a constant $y(r) = 1 - x(r)$ for all r.

A.4.2 Proof of Proposition 6

Denoting the Lagrange multiplier for the budget constraint by $\alpha \geq 0$, the Lagrangian can be written as

$$\alpha B_0 + \int_p^{\bar{r}} \left[-\frac{\hat{\Lambda}(r)}{f(r)} - \alpha \cdot \frac{[p - c - (1 - \delta) r] \hat{f}(r) - (1 - \delta) \hat{F}(r)}{f(r)} \right] y(r) \, \mathrm{d}F(r).$$

Observe that

$$-\frac{[p - c - (1 - \delta) r] \hat{f}(r) - (1 - \delta) \hat{F}(r)}{f(r)} = -\frac{(p - c) \hat{f}(r)}{f(r)} + (1 - \delta) \cdot \frac{r \hat{f}(r) + \hat{F}(r)}{f(r)}.$$

Here,

$$\frac{\mathrm{d}}{\mathrm{d}r} \left[\frac{\hat{f}(r)}{f(r)} \right] = \begin{cases} \dfrac{(1 - \delta) \left[-\frac{1-\delta}{\delta} f'\left(\frac{p-(1-\delta)r}{\delta} \right) f(r) - f'(r) f\left(\frac{p-(1-\delta)r}{\delta} \right) \right]}{\delta \left[f(r) \right]^2} & \text{for } p \leq r \leq \dfrac{p - \delta \underline{r}}{1 - \delta}, \\[20pt] 0 & \text{for } \dfrac{p - \delta \underline{r}}{1 - \delta} < r \leq \bar{r}. \end{cases}$$

Thus $\hat{f}(r)/f(r)$ increases on the interval $(p, (p - \delta \underline{r})/(1 - \delta))$, decreases discontinuously at the point $r = (p - \delta \underline{r})/(1 - \delta)$, and remains constant on the interval $((p - \delta \underline{r})/(1 - \delta), \bar{r})$. Moreover, by assumption,

$$(1 - \delta) r - (p - c) \Big|_{r = \frac{p - \delta \underline{r}}{1 - \delta}} = c - \delta \underline{r} > 0.$$

It follows that the objective function decreases discontinuously at the point $r = (p - \delta \underline{r})/(1 - \delta)$. Then rationing is required if

$$S(p) = 1 - F\left(\frac{p - \delta \underline{r}}{1 - \delta} \right).$$

Let p° be the solution to this equation; note that

$$S(p^\circ) = 1 - F\left(\frac{p^\circ - \delta \underline{r}}{1 - \delta} \right) \leq 1 - F(p^\circ) \implies p^\circ \leq p_c.$$

By continuity, rationing is required if p is in a neighborhood of p°.

A.4.3 Proof of Proposition 7

When the designer can intervene with a tax, consumers pay p for a unit of the private good but producers supply $S(p - \tau)$. Thus the design problem is

$$
\max_{x, \overline{U}} \ V(p - \tau) + \underline{U} \cdot \underline{\Lambda} + (1 - \delta) \int_p^{\overline{r}} \hat{\Lambda}(r) x(r) \ \mathrm{d}r + \int_p^{\overline{r}} (r - p) \ \mathrm{d}F(r)
$$

$$
\text{s.t.} \begin{cases} x : [p, \overline{r}] \to [0, 1] \text{ is decreasing; } \underline{U} \geq 0, \\[2mm] \displaystyle\int_p^{\overline{r}} [1 - x(r)] \ \mathrm{d}F(r) = S(p - \tau), \\[2mm] B + \tau \cdot S(p - \tau) - \underline{U} + \displaystyle\int_p^{\overline{r}} [p - c - (1 - \delta) \, r] \, x(r) \hat{f}(r) \ \mathrm{d}r - (1 - \delta) \int_p^{\overline{r}} \hat{F}(r) x(r) \ \mathrm{d}r \geq 0. \end{cases}
$$

The remainder of the proof is essentially identical to the proof of Theorem 1, and is hence omitted.

A.4.4 Proof of Proposition 8

Let $\alpha \geq 0$ and μ denote the Lagrange multipliers for the budget constraint and the market-clearing constraint respectively. The Lagrangian for the design problem is

$$
\mathcal{L}(x, \underline{U}, \mu, \alpha; p, \tau) = V(p - \tau) + \underline{U} \cdot \underline{\Lambda} + (1 - \delta) \int_p^{\overline{r}} \hat{\Lambda}(r) x(r) \ \mathrm{d}r + \int_p^{\overline{r}} (r - p) \ \mathrm{d}F(r)
$$

$$
+ \mu \left[\int_p^{\overline{r}} [1 - x(r)] \ \mathrm{d}F(r) - S(p - \tau) \right] + \alpha \left[B + \tau \cdot S(p - \tau) - \underline{U} \right]
$$

$$
+ \alpha \left[\int_p^{\overline{r}} [p - c - (1 - \delta) \, r] \, x(r) \hat{f}(r) \ \mathrm{d}r - (1 - \delta) \int_p^{\overline{r}} \hat{F}(r) x(r) \ \mathrm{d}r \right].
$$

By the envelope theorem (cf. Milgrom and Segal, 2002),

$$
\frac{\partial \mathcal{L}}{\partial p} = V'(p - \tau) - [1 - F(p)] - \mu \, [1 - x(p)] \, f(p) - \mu S'(p - \tau)
$$

$$
+ \alpha \tau S'(p - \tau) + \alpha \int_p^{\overline{r}} x(r) \hat{f}(r) \ \mathrm{d}r - \alpha \, (\delta p - c) \, x(p) \hat{f}(p).
$$

Suppose that c is sufficiently large, so that $\delta p^* - c < 0$ (a sufficient condition for this is $c > \delta \bar{r}$). Assume for the sake of contradiction that $1 \in \operatorname{im} x^*$, so that $x^*(p^*) = 1$. Then

$$\mu^* - \alpha^* \tau^* = \frac{V'(p^* - \tau^*) - [1 - F(p^*)] + \alpha^* \left[\int_{p^*}^{\bar{r}} x^*(r) \hat{f}(r) \, \mathrm{d}r - (\delta p^* - c) \, \hat{f}(p^*) \right]}{S'(p^* - \tau^*)} > 0.$$

Note that $\alpha^* \geq \underline{\Lambda}$. Thus the inequality above holds because $S' > 0$ and

$$V'(p^* - \tau^*) - [1 - F(p^*)] + \alpha^* \int_{p^*}^{\bar{r}} x^*(r) \hat{f}(r) \, \mathrm{d}r \geq S(p^* - \tau^*) - [1 - F(p^*)] + \underline{\Lambda} \int_{p^*}^{\bar{r}} x^*(r) \hat{f}(r) \, \mathrm{d}r$$

$$= \underline{\Lambda} \int_{p^*}^{\bar{r}} x^*(r) \hat{f}(r) \, \mathrm{d}r - \int_{p^*}^{\bar{r}} x^*(r) f(r) \, \mathrm{d}r \geq 0.$$

However, because $V'(p) \leq \underline{\Lambda} \cdot S(p)$ for every price p,

$$\left. \frac{\partial \mathcal{L}}{\partial \tau} \right|_{p^*, \tau^*} = -V'(p^* - \tau^*) + \mu^* S'(p^* - \tau^*) - \alpha^* \tau^* S'(p^* - \tau^*) + \alpha^* S(p^* - \tau^*) > 0.$$

This is a contradiction.

A.4.5 Proof of Proposition 9

Suppose on the contrary that $\tau^* = 0$. Then, by setting $\partial \mathcal{L}/\partial p$ to 0 for the optimal mechanism,

$$\mu^* = \frac{V'(p^*) - [1 - F(p^*)] + \alpha^* \left[\int_{p^*}^{\bar{r}} x^*(r) \hat{f}(r) \, \mathrm{d}r - (\delta p^* - c) \, x^*(p^*) \hat{f}(p^*) \right]}{[1 - x^*(p^*)] \, f(p^*) + S'(p^*)} > 0.$$

This inequality holds by the same argument given in the proof of Proposition 8. Then we have

$$\left. \frac{\partial \mathcal{L}}{\partial \tau} \right|_{p^*, \tau^*} = -V'(p^*) + \mu^* S'(p^*) + \alpha^* S(p^*) \geq \mu^* S'(p^*) > 0.$$

This shows that the designer can always strictly gain by setting a higher tax in the private market.

A.4.6 Proof of Proposition 10

This follows directly from the proof of Proposition 3.

A.4.7 Proof of Proposition 11

This follows straightforwardly from Theorems 3 and 4 of Dworczak Ⓡ Kominers Ⓡ Akbarpour (2021) and is omitted here.