

Optimal Redistribution Through Subsidies*

Zi Yang Kang[†] Mitchell Watt[‡]

This version: November 12, 2024

This paper is frequently updated; the latest version is available [here](#).

Abstract

In this paper, we develop a model of redistribution where a social planner, seeking to maximize weighted total surplus, can subsidize consumers who participate in a private market. We identify when subsidies can strictly improve upon the laissez-faire outcome, which depends on the correlation between consumers' demand and need. We characterize the optimal nonlinear subsidy by quantifying when—and for which units of the good—the social planner uses a full subsidy (i.e., free provision) rather than a partial subsidy or no subsidy. Our findings provide justifications for (i) free provision of a baseline quantity and (ii) subsidizing goods for which demand and need are positively correlated.

JEL classification: D47, D82, H21, H23, I38.

Keywords: subsidies, mechanism design, redistribution, topping up

* We are grateful to Paul Milgrom and Andy Skrzypacz for many illuminating discussions. We also thank Laura Doval, Piotr Dworzak, Liran Einav, Caroline Hoxby, Ravi Jagadeesan, Ilya Segal, Philipp Strack, Filip Tokarski, and Kai Hao Yang for their comments and suggestions. In addition, Kang gratefully acknowledges the support of the Center of Mathematical Sciences and Applications at Harvard University, and Watt gratefully acknowledges the support of the Koret Fellowship, the Ric Weiland Graduate Fellowship, and the Gale and Steve Kohlhagen Fellowship in Economics at Stanford University.

[†] Department of Economics, University of Toronto; zy.kang@utoronto.ca.

[‡] Department of Economics, Stanford University; mwatt@stanford.edu; (**Job Market Paper**).

1 Introduction

Suppose a social planner seeks to redistribute by subsidizing consumers participating in a private market for a good. Subsidies can be nonlinear: the social planner can adjust the marginal subsidy for each additional unit consumed. Subsidies are also costly, requiring the social planner to increase taxes or divert spending from other redistribution programs. When should the social planner offer any subsidy at all? If she does offer subsidies, how should she structure them to target consumers in need most effectively?

These questions are not merely theoretical. Globally, an estimated 1.5 billion people benefit from food subsidies (Alderman, Gentilini and Yemtsov, 2017). In both developed and developing countries, government agencies subsidize the consumption of a wide range of goods sold in private markets, including child care, education, energy, fuel, health care, and transportation. These subsidies are often nonlinear by design. For instance, the Supplemental Nutrition Assistance Program (SNAP) in the United States—commonly known as food stamps—provides a full subsidy for a fixed monetary amount of food but zero subsidies for additional purchases.

In this paper, we study the problem of optimal redistribution through consumption subsidies. In our model, consumers have access to a private market for a good, and a social planner can subsidize some or all units of the good purchased. We assume the social planner cannot tax private market consumption, for example, because subsidies are implemented by a government agency (such as the USDA designing food stamps) that lacks taxation powers. Drawing on recent literature on redistributive mechanism design, we model the social planner’s redistributive objective by assigning heterogeneous welfare weights to consumers. To focus on in-kind redistribution, we exclude lump-sum cash transfers within the mechanism, instead modeling the social planner’s opportunity cost of spending (which may depend on the availability of lump-sum transfers *outside* the mechanism) by assigning a welfare weight to subsidy expenditures.

The consumers’ access to the private market for the good constrains the social planner’s ability to redistribute. If the social planner could control the entire market, she might prefer to set high marginal prices for consumption levels chosen by low-need consumers. However, because consumers can always access the private market, the marginal price of any unit of consumption cannot exceed the private market price. Equivalently, the total subsidy a consumer receives must be an increasing function of the quantity consumed. This property of consumption subsidies limits the social planner’s ability to redistribute from low-need to high-need consumers. In quantity space, we show that this limit on the marginal price for each unit is equivalent to the requirement that each consumer’s total consumption exceeds his *laissez-faire* consumption level.

In our first main result, Theorem 1, we identify a sufficient statistic determining when the social planner strictly benefits from offering nonlinear subsidies. Specifically, the social planner can benefit from providing subsidies if and only if she can identify at least one consumer type for whom she would be willing to offer a cash transfer to *all* consumers with higher demand for the good than the identified type. Intuitively, if the social planner reimburses a small amount of spending for consumers who purchase more than the laissez-faire demand of the identified type, the increase in their weighted consumer surplus is linear in the reimbursement paid, while the costs of distorting consumption for a small group of lower types are second-order. As a result, the social planner strictly benefits from sufficiently small reimbursements, while the optimal nonlinear subsidy schedule can only improve upon that outcome. This argument relies on the fact that the social planner can use a nonlinear subsidy schedule: if she is restricted to setting a linear (ad valorem) subsidy, she must distort the consumption of *all* consumers upwards, making it more difficult to justify intervention in the market.

To better understand the constraints imposed by consumers' access to the private market, we contrast the implications of Theorem 1 in two benchmark cases: for goods with positive versus negative correlations between demand and need. When demand and need are *negatively correlated*—which might best model consumption patterns for food, education, and normal goods—the social planner can profitably intervene using nonlinear subsidies if and only if the *average* welfare weight exceeds the opportunity cost of funds. In this case, the social planner would prefer to make cash transfers to all consumers because total subsidy payments—which are increasing in quantity consumed—are strictly more regressive than the equivalent cash transfer.¹ When demand and need are *positively correlated*—which might best model consumption of staple foods, public transit, and inferior goods—the social planner can profitably intervene using nonlinear subsidies whenever the *maximum* welfare weight exceeds the opportunity cost of funds. In this case, the social planner can always redistribute to the highest-demand (and thus highest-need) consumers by offering a subsidy that distorts their consumption upwards without making it so generous as to attract low-demand (and thus low-need) consumers.

In our second main result, Theorem 2, we explicitly characterize the optimal subsidy mechanism. Whenever the average welfare weight exceeds the opportunity cost of public funds, the optimal subsidy mechanism includes free provision of a baseline quantity of the good. Conversely, whenever the average welfare weight is strictly less than the opportunity cost of public funds, the initial units of the good are not subsidized. In general, for any type for whom


¹ As we discuss below, there may be other reasons that the social planner prefers to use in-kind subsidies over cash transfers, including political preferences for in-kind redistribution and household dynamics.

the condition in Theorem 1 is satisfied, the *marginal* unit purchased by that type in the optimal subsidy mechanism is at least partially subsidized. The optimal subsidy schedule is *never* linear: the social planner can always improve over a linear subsidy mechanism by imposing a cap on the subsidy payment, a floor (or co-pay) on spending to qualify for subsidies, or by providing a quantity of the good for free.

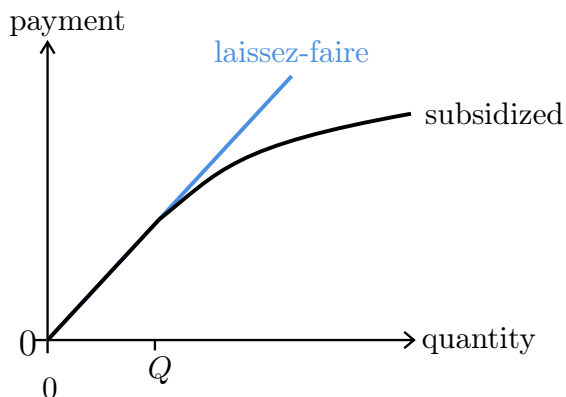
When demand and need are negatively correlated, Theorem 2 implies that it is optimal for the social planner to either offer a quantity of the good for free or no subsidies at all. Intuitively, this is because whenever the average welfare weight exceeds the opportunity cost of subsidy spending—which is the condition identified in Theorem 1 under which the social planner can benefit from subsidy intervention under negative correlation—the planner would like to offer cash transfers to all consumers. If the social planner can only redistribute in-kind, she uses free provision of the good as a substitute for that cash transfer. The optimal subsidized payment schedule with negative correlation generally contains at most three components, as illustrated in Figure 1(a): (i) free provision of some units of the good, (ii) partial subsidization of some additional units up to a maximum quantity, and (iii) any consumption beyond that maximum at the market price. In the negative correlation case, we find that, under the optimal mechanism, the social planner would strictly benefit from restricting subsidy recipients from purchasing additional units in the private market. With that restriction, the social planner could screen out low-need consumers while targeting subsidies to high-need ones.

On the other hand, when demand and need are positively correlated, we find that—under the optimal mechanism determined in Theorem 2—the social planner does not benefit from restricting subsidy recipients from purchasing additional units in the private market. The optimal mechanism, in that case, matches the optimal mechanism in a relaxed program in which the social planner can require consumers to choose between participating in a subsidy program or a private market (but not both)—a problem that we study in a companion paper (Kang and Watt, 2024). The resulting subsidy program has a *self-targeting* property: subsidies are provided precisely to the consumers that the social planner most wants to benefit, with the largest benefits accruing to those assigned the highest welfare weights. By screening consumers based on demand, in-kind subsidies can be a more effective instrument for redistribution than cash transfers. In that case—which is when the opportunity cost of funds is greater than the average welfare weight but less than the maximum welfare weight—the optimal subsidy involves some initial units of the good offered at the market price, with additional units beyond some threshold sold at a reduced price, as illustrated in Figure 1(b).

Our characterizations of optimal subsidy programs comport with the features of some existing



(a) Optimal payment schedule when welfare weights are negatively correlated with demand and, on average, higher than the opportunity cost of funds.



(b) Optimal payment schedule when welfare weights are correlated with demand and, on average, lower than the opportunity cost of funds.

Figure 1: Illustrating the optimal subsidized payment schedules.

programs. For instance, the Supplemental Nutrition Assistance Program (SNAP) provides a fixed monthly benefit with strict criteria for eligibility, and recipients are expected to supplement subsidized consumption with purchases in the private market. The resulting payment schedule for food mirrors the optimal mechanism when demand and need are negatively correlated, while its narrow eligibility may be a product of the need to restrict the program to a group of consumers with a high average welfare weight. In contrast, “fare capping” programs used by public transit authorities in various cities (including New York, London, Sydney and Hong Kong) cap total fares paid within a weekly or monthly period, allowing unlimited free trips once the cap is met, typically with broader eligibility. The payment schedule with fare caps resembles the optimal payment schedule when demand and need are positively correlated in that they both subsidize consumption beyond a minimum quantity consumed. The broad eligibility of consumers for fare-capping programs may also be due to the self-targeting nature of such mechanisms. Later in this paper, we explore the applicability and limitations of our results for other in-kind subsidy programs, such as those for pharmaceuticals and staple foods.

We also analyze how changes in economic primitives affect the optimal subsidy mechanism, focusing on three kinds of changes: an increase in the social planner’s preference for redistribution, a change in the correlation between demand and welfare weights, and an increase in demand for the good. Stronger redistributive preferences and increased demand for the good both lead to more generous subsidy programs. However, while stronger redistributive preferences expand the set of subsidized consumers, an increase in demand for the good leaves the set of subsidized consumers unchanged. The subsidy program is also more generous when there is a stronger correlation

between demand and need. As a result, holding all else equal, both the social planner and the average eligible consumer favor programs that are restricted to consumers with higher welfare weights and subsidize goods with a strong positive association between demand and need.

Our results may help explain real-world examples of subsidy programs that have strengthened eligibility requirements or reduced product scope over time. For example, SNAP’s eligibility requirements have been tightened multiple times to better target the program to more vulnerable populations. In the Egyptian Bread Subsidy Program, the government limits the weight of the loaves and the quality of wheat that may be used in the production of subsidized bread. The Indonesian government has recently restricted access to its fuel subsidy program to ride-share operators and owners of vehicles with smaller engines. Each of these reforms may have increased the average welfare weight of eligible consumers or the correlation between consumption and the social planner’s preferences for redistribution. We discuss these examples further in Section 7 below.

We also discuss extensions of our model to address two other important considerations for subsidy design: the equilibrium effects of subsidies on market prices and how the subsidy design problem changes when the social planner has some ability to tax the outside market for the good. These extensions enable our findings to apply to markets where the social planner can influence the prices that consumers face for supplemental purchases. We also describe how the social planner’s welfare weights—assumed to be exogenous in our baseline model—can be endogenized in a richer model without qualitatively changing our main results.

Our main technical contribution in this paper is an explicit characterization of the solution of mechanism design problems with lower-bound constraints on the implementable allocation rules. In our subsidy design problem, this constraint arises from the consumers’ ability to purchase units at the private market price, which ensures that no consumer can receive less than his *laissez-faire* consumption. Handling this constraint is challenging because it interacts with the standard monotonicity constraint, necessitating a solution method that incorporates both constraints simultaneously. In particular, the optimum is not obtained by solving a relaxed problem without the lower-bound constraint and enforcing the constraint on its solution. Instead, we introduce a new “double ironing” procedure accounting for these interactions. We also deviate from the standard mechanism design approach to verify the optimality of the allocation rule. As is standard in mechanism design, we identify an unconstrained convex program for which the allocation rule is a pointwise maximizer. But instead of using a Lagrangian approach, we use variational inequality methods to relate the optimality conditions of the two convex programs.

1.1 Related Literature

Our paper is related to the public finance literature on optimal commodity taxation (and subsidization), which has focused on designing tax systems that maximize social welfare while balancing government revenue needs. [Ramsey \(1927\)](#) initiated this field with the “inverse elasticity rule,” suggesting that goods with lower demand elasticities should be taxed more heavily to minimize distortions. [Diamond \(1975\)](#) incorporated equity considerations by generalizing [Ramsey’s](#) rule to multiperson economies, finding that optimal taxes should reflect not only the elasticity of demand but also distributional concerns. Like [Diamond](#), we find that subsidies should be focused on goods with a positive correlation between demand and need; however, unlike [Diamond](#), we allow for nonlinear subsidy rules, explicitly characterizing when the social planner benefits from using subsidies and showing that nonlinearity increases the scope of redistribution through subsidies. While the literature has developed explicit expressions for optimal nonlinear income tax schedules (e.g., [Mirrlees, 1976](#); [Stiglitz, 1987](#)), explicit solutions for nonlinear commodity tax schedules are generally more difficult to obtain. [Mirrlees \(1986\)](#) derived first-order conditions for optimal nonlinear commodity taxes, yet using those results to identify properties of the optimal tax schedule is challenging, partly because they are stated implicitly in terms of optimal Lagrange multipliers that are not computed. In contrast, our mechanism design approach determines an explicit characterization of the optimal allocation rule and subsidy schedule (effectively calculating the optimal Lagrange multipliers), which allows us to derive additional properties and conduct comparative statics. Our approach also differs from traditional optimal taxation models by focusing on subsidization in settings where the social planner cannot impose taxes; we show that the optimal subsidy rule is not simply the optimal tax-and-subsidy rule with taxes set to zero.

On the other hand, the celebrated theorem of [Atkinson and Stiglitz \(1976\)](#) questions the use of differential commodity taxation (and subsidization) altogether, finding that—under a separability assumption on consumption and labor preferences—an optimal nonlinear income tax alone is sufficient for redistribution. The [Atkinson and Stiglitz](#) theorem does not apply in our setting due to a failure of “incentive-separability” in our model ([Doligalski, Dworzak, Krysta and Tokarski, 2023](#); see also related discussions in [Saez, 2002](#), and [Pai and Strack, 2024](#)). Intuitively, because the consumer’s private information about their consumption preferences is informative about the social planner’s welfare weight, the social planner can use commodity subsidization to elicit private information relevant to determining the recipient’s need for redistribution.

Various papers in public finance have focused explicitly on in-kind transfers (including many

surveyed by [Currie and Gahvari, 2008](#)). We contribute to that literature by deriving the optimal form of in-kind transfers (namely, direct provision or partial subsidization) as a function of the demand and supply of the good and the social planner’s objective. Stemming from the pioneering work of [Nichols and Zeckhauser \(1982\)](#), this literature has shown how in-kind transfers can screen individuals better than cash transfers (e.g., [Blackorby and Donaldson, 1988](#)) and how a private market might affect this screening (e.g., [Besley and Coate, 1991](#); [Gahvari and Mattos, 2007](#)). Within this literature, the closest paper is [Blomquist and Christiansen \(1998\)](#), which focuses on the tradeoff between allowing subsidy recipients to top up their subsidized allocations and restricting topping up in a subsidy program using a Mirrleesian model of labor and leisure. Whereas previous work has tended to assume the form of in-kind transfer that the social planner can use—typically an *ad valorem* subsidy—our paper endogenizes the social planner’s choice using tools from mechanism design and finds that the optimal subsidy mechanism is nonlinear. We also focus on screening using consumption preferences, and only briefly discuss costly screening (“ordeals”), as studied by [Nichols and Zeckhauser \(1982\)](#), [Finkelstein and Notowidigdo \(2019\)](#), [Yang \(2021\)](#), [Yang, Dworzak and Akbarpour \(2024\)](#) and others.

Our paper is also related to the growing literature on redistributive mechanism design. [Weitzman \(1977\)](#) first observed that distortions from the allocation that would arise in a competitive market can help redistribute when the social planner seeks to maximize a different objective than utilitarian welfare. An ensuing literature has used tools from mechanism design to formalize [Weitzman’s](#) observation and characterize optimal mechanisms in general settings (e.g., [Condorelli, 2013](#); [Che, Gale and Kim, 2013](#); [Dworczak !\[\]\(c8d96c8885d3000a912c2582004aed63_img.jpg\) Kominers !\[\]\(3ad821e3ca7dd4cb7003e9c8d982e254_img.jpg\) Akbarpour, 2021](#); [Akbarpour !\[\]\(177bde115c7ebbeffa559d05eea9e94b_img.jpg\) Dworzak !\[\]\(cab2e95699b614c49dd80341e1932607_img.jpg\) Kominers, 2024b](#)). These insights have also been applied to settings with finitely many agents (e.g., [Kang and Zheng, 2020](#); [Reuter and Groh, 2020](#)) and externalities (e.g., [Kang, 2024](#); [Akbarpour !\[\]\(75f9a43febaa9aa08b77b73c8ad8a855_img.jpg\) Budish !\[\]\(cc11c320e6649662ebbe761994792de4_img.jpg\) Dworzak !\[\]\(f8f94fe68788e4fac99840bca8324d12_img.jpg\) Kominers, 2024a](#); [Pai and Strack, 2024](#)). In this literature, our paper is closest to our companion paper, [Kang and Watt \(2024\)](#), which studies the design of in-kind redistribution programs when consumers must choose between participating in a subsidized program or the private market (as in public housing subsidy programs and some others). Our paper is also related to [Kang \(2023\)](#), in that we allow consumers to access a private market that the social planner cannot design. While [Kang \(2023\)](#) assumes that the social planner is restricted to only some forms of in-kind transfers (namely, only a good of a single quality level can be directly provided) and focuses on equilibrium effects when consumers must choose between subsidized public consumption or private market consumption, our paper contributes by removing this restriction and endogenizing the social planner’s choice of in-kind transfers.

The presence of a private market that the social planner cannot control in our model connects our paper to work on partial mechanism design, or “mechanism design with a competitive fringe.” This literature has studied optimal interventions in markets with adverse selection (e.g., [Philippon and Skreta, 2012](#); [Tirole, 2012](#); [Fuchs and Skrzypacz, 2015](#)), optimal pricing with resale (e.g., [Carroll and Segal, 2019](#); [Dworczak, 2020](#); [Loertscher and Muir, 2022](#)), optimal contracting with type-dependent outside options (e.g., [Jullien, 2000](#)), optimal contracting between firms (e.g., [Calzolari and Denicolò, 2015](#); [Kang and Muir, 2022](#)), and optimal redistribution (e.g., [Kang, 2023](#)). Our paper contributes to this literature by enriching how the social planner can intervene: whereas previous work has tended to assume that the social planner and the private market use different production technologies for tractability, our paper develops tools to analyze when the social planner and the private market can access the same production technology (which we interpret as the ability of the social planner to directly subsidize private purchases or to costlessly contract public production out to private firms).

From a methodological perspective, our paper primarily builds on two different techniques in the literature on mechanism design. On one hand, we use insights from [Kang \(2024\)](#) in formulating our mechanism design problem as a projection problem: the social planner seeks to project a “target function” onto the set of feasible allocation functions. The projection operator is more complicated in this paper because the private market constrains the set of feasible allocation functions, but these insights nonetheless allow for an explicit formula for the optimal allocation function, allowing us to study comparative statics of the optimal mechanism. On the other hand, using a proof technique to similar to [Toikka \(2011\)](#), we relate the optimality conditions for our constrained program to the optimality conditions of a related unconstrained convex program, for which the solution is a pointwise maximizer. We show how to adapt the ironing operation introduced by [Myerson \(1981\)](#) (and extended by [Toikka, 2011](#) and others) to account for the private market constraint. We verify our candidate solution using variational inequality methods (distinct from the Lagrangian techniques developed for similar problems by [Amador, Werning and Angeletos, 2006](#) and [Amador and Bagwell, 2013](#)).

Finally, as we show, the constraints on the set of feasible allocations due to the private market lead to a lower-bound constraint on the allocation rule, an example of what the literature has called “first-order stochastic dominance constraints.” This connects our paper to a few papers that have studied mechanism design problems under such constraints (e.g., [Kang and Muir, 2022](#); [Corrao, Flynn and Sastry, 2023](#); [Yang and Zentefis, 2024](#)). Our paper contributes both in the application we study and by obtaining an *explicit* formula for the optimal allocation function for problems with such constraints for separable (cf. [Mussa and Rosen, 1978](#)) preferences.

1.2 Organization

In Section 2, we introduce a model of subsidy design for a redistributive social planner in a setting where consumers have access to a private market. In Section 3, we identify when the social planner benefits from offering subsidies. In Section 4, we characterize the optimal subsidy mechanism and discuss the details of its implementation. In Section 5, we discuss how various economic primitives affect the optimal subsidy mechanism, including the differences in subsidy design according to the correlation between demand and need, as well as the effect of changes in the designer’s preferences for redistribution. In Section 6, we discuss various extensions of the model, and in Section 7, we discuss how our findings relate to in-kind subsidy designs observed in practice and how our results provide insights into product choice and eligibility restrictions for in-kind subsidies, which are not explicitly included in our baseline model. In Section 8, we conclude. Our appendices contain proofs omitted from the main text and additional material.

2 Model

In this section, we develop a model of redistribution through nonlinear subsidies. We start with a standard model of a private market and describe the laissez-faire equilibrium. We then formulate the social planner’s problem and conclude with a discussion of modeling assumptions.

2.1 Private Market

There is a unit mass of risk-neutral consumers in the market for a homogeneous, divisible good. The good is supplied competitively by producers in a private market. For simplicity, we focus on the case where producers face a constant marginal cost $c > 0$ for each unit of the good; we show how our analysis extends to the case of a nonconstant marginal cost (i.e., an upward-sloping supply curve) in Section 6.1.

While consumers have quasilinear preferences over money, they differ in their preferences over consumption quantity. We model heterogeneity in consumption preferences by identifying each consumer with a type $\theta \in [\underline{\theta}, \bar{\theta}] =: \Theta \subset \mathbb{R}_{++}$. A consumer of type θ derives a utility of $\theta v(q) - t$ from consuming q units of the good at a total price of t . Each consumer is privately informed about his type. Types are drawn from a cumulative distribution function F , which is absolutely continuous and has a positive density function f supported on Θ .

Consumers have diminishing marginal utility for consumption. To capture this, we assume

that $q \in [0, A]$, where A is the maximum quantity of the good demanded by a consumer at any price, and that the function $v : [0, A] \rightarrow \mathbb{R}_+$ is twice continuously differentiable, with $v' > 0$ and $v'' < 0$ (i.e., v is increasing and strictly concave).² For notational simplicity, we extend the domain of $(v')^{-1}$ to \mathbb{R} by setting $(v')^{-1}(z) = 0$ for $z \geq v'(0)$ and $(v')^{-1}(z) = A$ for $z \leq v'(A)$. This allows individual demand curves to be written as

$$D(p, \theta) = (v')^{-1} \left(\frac{p}{\theta} \right).$$

2.2 Laissez-Faire Equilibrium

Next, we describe the laissez-faire equilibrium. Because the market is perfectly competitive, each unit of the good is priced at its marginal cost, c . Given that price, each consumer solves his utility maximization problem,

$$U^{\text{LF}}(\theta) := \max_{q \in [0, A]} [\theta v(q) - cq].$$

Let $q^{\text{LF}}(\theta) = D(c, \theta)$ denote the solution to the consumer's utility maximization problem (i.e., the laissez-faire allocation function); since v is strictly concave, $q^{\text{LF}}(\theta)$ is uniquely defined.

2.3 Social Planner's Problem

We now study the problem faced by a social planner who wishes to redistribute through subsidies. The planner assigns weights to the utility of different consumers and the costs of spending as follows:³

- (a) *Consumer Surplus*: The social planner assigns a positive *welfare weight* $\omega(\theta)$ to type θ 's utility, corresponding to the increase in expected social welfare associated with giving a dollar to a consumer of type θ .⁴ We assume that $\omega : \Theta \rightarrow \mathbb{R}_+$ is continuous.

While our main results (Theorems 1 and 2) do not require additional assumptions on ω , we focus on two key cases to study the economic implications:

² As we show in Appendix D.3, our analysis extends to the case where consumers have constant marginal utility for consumption.

³ In Section 6.3, we discuss how these weights may be determined endogenously in a richer model with budget constraints and concave utility, with the weight α corresponding to the Lagrange multiplier on the social planner's budget constraint and $\omega(\theta)$ as the expected marginal value for money of a consumer with type θ .

⁴ While the welfare weight is a deterministic function of type in our model, it is equivalent to model the consumer as having a type $(\theta, \tilde{\omega})$ where $\tilde{\omega}$, the welfare weights, is a random variable, in which case $\omega(\theta) = \mathbf{E}[\tilde{\omega}|\theta]$ is the expected welfare weight of a consumer with demand type θ (see, e.g., Akbarpour [\(r\)](#) al. (2024b)). In Section 7, we will return to this interpretation of $\omega(\theta)$.

- (i) *Negative correlation*: ω decreases in θ , meaning higher-demand consumers have lower welfare weights. If ω is determined by income, with the social planner seeking to redistribute to lower-income consumers, the negative correlation assumption may best model markets for food, housing, education, and other normal goods.
 - (ii) *Positive correlation*: ω increases in θ , meaning higher-demand consumers have higher welfare weights. If ω is determined by income, with the social planner seeking to redistribute to lower-income consumers, the positive correlation assumption may best model markets for staple food, public transportation, and inferior goods.
- (b) *Cost of Spending*: The social planner assigns a weight α to the cost of spending, capturing the *opportunity cost of funds*. If $\omega(\theta)$ exceeds α , the social planner experiences a net benefit from transferring a dollar to a consumer of type θ ; conversely, if α exceeds $\omega(\theta)$, the social planner would benefit from receiving a dollar from a consumer of type θ .

Below, we show that the key features of the optimal subsidy design depend on how α compares to $\mathbf{E}[\omega]$: we say that the opportunity cost of funds is *high* if $\alpha > \mathbf{E}[\omega]$ and that the opportunity cost of funds is *low* if $\alpha \leq \mathbf{E}[\omega]$. The opportunity cost of funds might be high because the social planner has other redistributive programs competing for the same budget (including cash transfers outside the mechanism) or because of distortions associated with raising tax revenue. On the other hand, we would expect α to be low if the social planner faces political constraints that limit her ability to redistribute outside the mechanism.⁵

The social planner can intervene in the private market by offering nonlinear subsidies. As explained in Section 1, our motivation to consider nonlinear subsidies stems from the fact that, in practice, many subsidy programs are nonlinear.

To model nonlinear subsidies, we allow the social planner to specify a marginal subsidy for each incremental unit of consumption. We require the marginal subsidy for each unit of consumption to be nonnegative, recognizing that each consumer can always purchase a unit of the good directly from the private market without accessing the subsidy schedule. Formally, the social planner chooses a marginal subsidy function $\sigma : [0, A] \rightarrow \mathbb{R}_+$, where $\sigma(q)$ denotes the marginal subsidy for the q^{th} unit of consumption. Given the chosen marginal subsidies, each consumer faces the net

⁵ For example, [Liscow and Pershing \(2022\)](#) find that the U.S. general population largely prefers in-kind subsidies to cash transfers. In other cases, subsidies may be designed by government agencies without the authority to offer cash transfers. Alternatively, in-kind subsidies might be preferred for household economics reasons: for example, [Currie \(1994\)](#) finds evidence that in-kind subsidy programs may have stronger benefits for children than cash transfer programs.

price schedule $P^\sigma(q) = cq - \Sigma(q)$, where

$$\Sigma(q) := \int_0^q \sigma(z) \, dz,$$

is the total subsidy paid for q units of consumption.

For example, the social planner may choose to use a linear subsidy by having a constant marginal subsidy $\sigma(q) \equiv \hat{\sigma}$; this yields a net (linear) price schedule of $P^\sigma(q) = (c - \hat{\sigma})q$. However, the social planner may alternatively specify a nonconstant marginal subsidy function. For instance, she may choose to set $\sigma(q) = c \cdot \mathbf{1}_{q \leq \underline{q}}$, which yields a full subsidy on the first \underline{q} units of consumption and no subsidy thereafter.

To distinguish between subsidies for the good and cash transfers, we restrict the social planner to choose only subsidies that do not exceed the cost of the good:

$$\Sigma(q) \leq cq \quad \text{for any } q \in [0, A]. \quad (\text{SUB})$$

This condition ensures that consumers never face a negative net price for any quantity of the good. Equivalently, this condition requires that consumer payments are always bounded below by zero; hence, the social planner can only subsidize consumption but not give consumers cash.

Next, to simplify our analysis, we equivalently reformulate the social planner's problem as one of choosing the induced allocation received by each consumer and the corresponding payment. This reformulation enables us to apply mechanism design tools. Specifically, the social planner chooses a direct mechanism (q, t) , which consists of:

- (i) an allocation function $q : \Theta \rightarrow [0, A]$, where $q(\theta)$ is the quantity consumed by a consumer with type θ ; and
- (ii) a payment function $t : \Theta \rightarrow \mathbb{R}$, where $t(\theta)$ is the payment, net of subsidies, made by a consumer with type θ .

We say that a mechanism (q, t) is *implemented* by a marginal subsidy function $\sigma : [0, A] \rightarrow \mathbb{R}_+$ if σ satisfies (SUB) and

$$q(\theta) \in \arg \max_{z \in [0, A]} [\theta v(z) - P^\sigma(z)], \text{ and } t(\theta) = P^\sigma(q(\theta)).$$

We now describe the feasibility constraints on mechanisms that may be implemented using subsidies by the social planner.

First, by the revelation principle (Myerson, 1981), the social planner can restrict attention to incentive-compatible mechanisms without loss of generality, so that consumers can do no better than to truthfully report their types:

$$\theta \in \arg \max_{\theta' \in \Theta} [\theta v(q(\theta')) - t(\theta')] \quad \text{for all } \theta \in \Theta. \quad (\text{IC})$$

Second, unlike standard mechanism design problems, the fact that the mechanism is induced by a nonlinear subsidy results in a lower-bound constraint on the allocation function:

Lemma 1. *Given any allocation rule $q : [0, A] \rightarrow \mathbb{R}_+$, there exists a mechanism (q, t) satisfying (IC) and a marginal subsidy function $\sigma : [0, A] \rightarrow \mathbb{R}_+$ implementing (q, t) if and only if q is nondecreasing in θ and satisfies*

$$q(\theta) \geq q^{\text{LF}}(\theta) \quad \text{for all } \theta \in \Theta. \quad (\text{LB})$$

Intuitively, subsidies incentivize increased consumption. As Lemma 1 shows, the converse holds as well: any allocation rule with increased consumption can be implemented by nonlinear subsidies. The proof of Lemma 1 is in Appendix A.

Third, our focus on consumption subsidies leads to a constraint on payments within the mechanism and on the feasible payoffs for consumers:

Lemma 2. *Let (q, t) be a mechanism that satisfies (IC) and (LB). Then there exists a marginal subsidy function $\sigma : [0, A] \rightarrow \mathbb{R}_+$ implementing (q, t) if and only if*

$$t(\theta) \geq 0 \quad \text{for all } \theta \in \Theta, \quad (\text{NLS})$$

and $\underline{U} := \theta v(q(\underline{\theta})) - t(\underline{\theta})$ satisfies

$$\underline{U} \geq U^{\text{LF}}(\underline{\theta}). \quad (\text{IR}_{\underline{\theta}})$$

The (NLS) constraint means that while the social planner can subsidize goods fully (i.e., setting a price of zero for any amount of the good below some threshold quantity), she cannot give consumers money. However, we allow for cash transfers outside of the mechanism, which we model through the social planner's preferences over subsidy spending, captured by the weight α . Moreover, the fact that consumers are always subsidized (and not taxed) implies that the utility of each agent must weakly exceed their laissez-faire utility; Lemma 2 clarifies that it suffices to enforce that individual rationality constraint for the lowest type only, with (IC) and (LB) implying the condition for all higher types. The proof of Lemma 2 is in Appendix A.

Putting these constraints together, the social planner chooses a feasible mechanism to maximize weighted total surplus:

$$\max_{(q,t)} \int_{\Theta} \left\{ \omega(\theta) \underbrace{[\theta v(q(\theta)) - t(\theta)]}_{\text{consumer surplus}} - \alpha \underbrace{[cq(\theta) - t(\theta)]}_{\text{total costs}} \right\} dF(\theta), \quad (\text{OPT})$$

such that (q, t) satisfies (IC), (LB), (NLS), and (IR $_{\theta}$).

Because the laissez-faire mechanism $(q^{\text{LF}}, cq^{\text{LF}})$ is always feasible, the set of feasible mechanisms is always nonempty, and the optimization program (OPT) is well-posed. In Appendix C.2, we show that there exists a unique optimal mechanism when $\mathbf{E}[\omega] \neq \alpha$, and when $\mathbf{E}[\omega] = \alpha$, the optimal mechanism is unique up to a constant in the payment rule. We let (q^*, t^*) denote an optimal subsidy mechanism, and call q^* the *subsidy allocation rule*. We also write $U^*(\theta)$ for the consumer surplus of type θ in the optimal subsidy mechanism.

2.4 Model Discussion

The key assumption in our model is that the social planner *subsidizes* but does not *tax* private market consumption of the good. Without this assumption, the social planner could set an infinite tax on the private market and implement any nondecreasing allocation rule using appropriate subsidies. Ruling out taxes is realistic in many cases, such as when the subsidy designer lacks taxation powers (e.g., because the subsidy is designed by a local government authority or agency separate from the taxation office), when there are political constraints that make commodity taxation costly, or when the social planner internalizes the costs of high taxation on consumers ineligible for subsidies. In other markets, while it may be possible to tax, restrict, or shut down the private market, such actions come at a cost, which the social planner would need to weigh against the benefits of better-targeted subsidies. The optimal subsidy program identified in this paper helps assess the tradeoffs of such interventions, which we discuss further in Section 5.4.

An equivalent interpretation of our model is that the social planner produces quantities of the good using the same technology as the private market and sells the units it produces to eligible consumers at a subsidized price schedule. In that case, the social planner's inability to tax units of the good is equivalent to assuming that consumers can always *top up* their subsidized consumption in the private market for the good, as is the case in many real-world subsidy markets (Currie and Gahvari, 2008). In many markets, the prohibitive costs associated with monitoring and enforcing private market restrictions makes topping up practically unavoidable. In other

markets, our characterization of the optimal subsidy program helps assess the tradeoffs associated with enforcing restrictions on private market consumption, as we discuss further in Section 6.2.

Another important assumption in our subsidy design model is the no lump-sum transfers constraint, (NLS). In practice, the inability to make lump-sum transfers may arise from political constraints, administrative costs, or the fact that the subsidy program is designed by an agency (e.g., the transportation department) without the authority to make unrestricted cash transfers. While we rule out lump-sum transfers directly within the mechanism, any such transfers made outside the mechanism can be captured in our model via their influence on the opportunity cost of funds α . Including the (NLS) constraint allows us to assess when the social planner would prefer cash transfers (alongside or in place of a subsidy) and when the restriction is non-binding.

We have made two significant assumptions on production. The first is that the social planner can access the same production technology as the private market, which may reflect arrangements with either producers or consumers.⁶ On the supply side, for example, the social planner could costlessly contract with private firms to produce subsidized goods or subsidize them to offer the recommended pricing schedules. On the demand side, the social planner could reimburse consumers for their private market spending. The second is that the supply schedule is perfectly elastic. That assumption shuts down the possibility of equilibrium effects on prices caused by the social planner's choice of subsidy mechanism. We foreclose that possibility in our main model to concentrate on the effects of the consumers' access to the private market, with other papers (including Kang, 2023) focusing on equilibrium effects in isolation. However, in Section 6.1, we discuss how to extend our baseline model to incorporate equilibrium effects on prices.

In our formulation, we have also implicitly assumed that all consumers are eligible for the subsidy program and that the choice of product for subsidization is exogenous. In practice, the social planner may have some discretion over which product to subsidize and which consumers qualify based on observable characteristics, a possibility we revisit in Section 7.

3 When Does The Social Planner Offer Subsidies?

In this section, we present our first main result: a necessary and sufficient condition for when the social planner can strictly improve on the laissez-faire outcome using subsidies.

⁶ By contrast, Kang (2023) studies the case in which private producers are more efficient than the social planner who can offer only one price-quantity pair to consumers (who cannot top up in the private market).

Theorem 1 (scope of optimal subsidy). *The optimal mechanism (q^*, t^*) strictly improves on the laissez-faire outcome if and only if*

$$\max_{\hat{\theta} \in \Theta} \mathbf{E}_\theta \left[\omega(\theta) \mid \theta \geq \hat{\theta} \right] > \alpha.$$

In particular:

- (a) *If ω is decreasing in θ , then (q^*, t^*) strictly improves on the laissez-faire outcome if and only if $\mathbf{E}_\theta[\omega] > \alpha$.*
- (b) *If ω is increasing in θ , then (q^*, t^*) strictly improves on the laissez-faire outcome if and only if $\max \omega > \alpha$.*

Theorem 1 provides a sufficient statistic that determines when the social planner can benefit from intervention in the market using nonlinear subsidies. It illustrates that there is a greater scope for redistribution through nonlinear subsidies when demand and need are positively correlated than when they are negatively correlated, extending similar results for linear subsidies known in the public finance literature (e.g., [Atkinson and Stiglitz, 2015](#); [Mackenzie, 1991](#)). However, as we discuss further below, there is a greater scope for redistribution using nonlinear subsidies than when the social planner is restricted to choosing linear subsidy instruments.

3.1 Proof of Theorem 1

Before turning to the proof of Theorem 1, we note a key property of consumption subsidies driving the result.

Lemma 3. *Suppose that (q, t) is implemented by marginal subsidy function $\sigma : [0, A] \rightarrow \mathbb{R}_+$. Then the total subsidies received by a consumer of type θ , $\Sigma(q(\theta))$, is a nondecreasing function of θ .*

Proof. Note that total subsidies $\Sigma(q)$ is a nondecreasing function of q because the marginal subsidy $\sigma(q) \geq 0$. On the other hand, because $q(\theta)$ is the maximizer of $\theta v(z) - P^\sigma(z)$, which is supermodular, $q(\theta)$ is nondecreasing in θ by the Topkis Theorem (cf. [Milgrom and Shannon, 1994](#)). This implies that $\Sigma(q(\theta))$ is the composition of two nondecreasing functions, so it is nondecreasing in θ as well. \square

The requirement that total subsidies increase with consumer type means that the social planner's redistributive power depends critically on the relationship between welfare weights and

consumer demand. Broadly, when welfare weights and demand are negatively correlated—as might be relevant for food, education, and other normal goods—the requirement that subsidies increase with type is in conflict with the social planner’s redistributive objective. On the other hand, when welfare weights and demand are positively correlated—as might be relevant for staple foods, public transit, and inferior goods—the monotonicity of subsidies is aligned with the social planner’s redistributive objective.

We now turn to the proof of Theorem 1.

Proof. We first prove the “only if” direction of Theorem 1. To establish this result, we must show that the social planner does not benefit from intervention when the condition described in Theorem 1 fails. We focus here on the proof of this result in the case that ω is decreasing in θ which illustrates the main intuition; the generalized proof is presented in Appendix C.3.

Note that a consumer’s increase in surplus from a subsidy is bounded above by the subsidy itself. This is because the consumer may need to spend part of the subsidy on extra consumption to qualify for the payment, and the value placed on consumption beyond the laissez-faire level is always less than the cost.⁷

When ω is decreasing, the social planner’s benefit of a marginal subsidy function σ , measured by the expected increase in weighted consumer surplus, is bounded as follows:

$$\begin{aligned} \int_{\Theta} [\omega(\theta)(U(\theta) - U^{\text{LF}}(\theta))] \, dF(\theta) &\leq \int_{\Theta} [\omega(\theta)\Sigma(q(\theta))] \, dF(\theta) \\ &\leq \mathbf{E}_{\theta}[\omega(\theta)] \mathbf{E}_{\theta}[\Sigma(q(\theta))], \end{aligned}$$

where the first inequality reflects the above reasoning, and the second arises from the negative correlation between ω (decreasing in θ by assumption) and the subsidy (increasing in θ by Lemma 3). This shows that subsidies are more regressive than the equivalent lump-sum transfer, so the social planner’s benefit is lower than that of an equivalent lump-sum transfer to all consumers.⁸ Whenever $\mathbf{E}[\omega] \leq \alpha$, the social planner prefers not to make a lump-sum transfer to all consumers, and, as a result, the social planner does not benefit from offering the more regressive subsidy program either.

⁷ Formally, because $\theta v(q(\theta)) - cq(\theta) \leq \theta v(q^{\text{LF}}(\theta)) - cq^{\text{LF}}(\theta)$ by definition of q^{LF} , we have $U(\theta) - U^{\text{LF}}(\theta) = \theta v(q(\theta)) - t(\theta) - [\theta v(q^{\text{LF}}(\theta)) - cq^{\text{LF}}(\theta)] \leq cq(\theta) - t(\theta) = \Sigma(q(\theta))$.

⁸ The fact that linear subsidies are more regressive than cash transfers in the case of negative correlation is known (see, e.g., Diamond and Mirrlees (1972)), but our result is stronger, implying that *any* subsidy offered to consumers with unrestricted access to the private market for the good must be more regressive than cash transfers.




Figure 2: The ε -perturbed subsidy schedule, illustrating an indifference curve of type $\tilde{\theta} \leq \hat{\theta}$ just indifferent between $q^{\text{LF}}(\tilde{\theta})$ and the subsidized $q^{\text{LF}}(\hat{\theta})$. Types between $\tilde{\theta}$ and $\hat{\theta}$ strictly prefer the subsidized $q^{\text{LF}}(\hat{\theta})$.

We now prove the “if” direction of Theorem 1 by showing how to construct a subsidized payment schedule that increases the social planner’s payoff compared to laissez-faire when $\mathbf{E}_{\theta}[\omega(\theta)|\theta \geq \hat{\theta}] > \alpha$ for some $\hat{\theta} \in \Theta$. In particular, consider the subsidized payment schedule illustrated in Figure 2, where the social planner reimburses ε of spending for any consumer buying more than $q^{\text{LF}}(\hat{\theta})$ units. All consumers with types $\theta \geq \hat{\theta}$ receive an ε increase in utility, while the social planner incurs a cost of ε . The overall effect on weighted total surplus is thus ε times the positive $\int_{\hat{\theta}}^{\bar{\theta}} [\omega(\theta) - \alpha] dF(\theta)$, which is linear in ε . On the other hand, the (IC) constraint requires the social planner to increase consumption for consumers just below $\hat{\theta}$ to $q^{\text{LF}}(\hat{\theta})$, as shown in Figure 2. While this raises their utility, the net effect after accounting for the cost of funds may be negative (but bounded below by $-\alpha\varepsilon$ per consumer). Because the mass of such consumers is $O(\sqrt{\varepsilon})$, which is decreasing in ε ,⁹ the total cost of the subsidy to those types is $O(\varepsilon^{\frac{3}{2}})$, which is negligible compared to the benefits of subsidizing those with $\theta \geq \hat{\theta}$ for small ε (which is linear in ε). Thus, the social planner can always design a subsidy scheme that increases its objective compared to the laissez-faire outcome. \square

3.2 Discussion

To further understand the limitations on redistribution facing the social planner in our model as a consequence of the consumers’ access the private market, we now compare Theorem 1 to similar results obtained for alternative models of government intervention.

⁹ We prove this formally in Appendix C.3.

Shutdown Benchmark. Suppose that the social planner can *shut down* the private market and become the sole producer and seller of the good. In that case, the social planner can set the marginal price of any unit of consumption to be higher than c , allowing her to extract revenue from consumers with higher welfare weights and use it to subsidize consumers with lower welfare weights.

To illustrate why this ability to tax may widen the scope of intervention, suppose ω is decreasing in θ and $\mathbf{E}[\omega] \leq \alpha$. In that case, Theorem 1 implies that a social planner who can only subsidize private consumption does not benefit from intervening in the private market. However, a social planner who can tax consumption of the good can improve over the laissez-faire outcome by taxing all consumption of the good. In particular, because $\mathbf{E}[\omega] \leq \alpha$, the social planner benefits from taxing *all* consumers, with higher taxes paid by consumers who consume higher quantities of the good. In that case (by an argument similar to the one in the proof of Theorem 1), the negative correlation between total tax payments and welfare weights implies that the benefit of the nonlinear taxation scheme *exceeds* the benefit of a lump-sum tax. In general, we have the following characterization of the scope of redistribution when the social planner can control the entire market for the good:

Proposition 1 (scope of intervention with private market shutdown). *Suppose the social planner is the sole producer and seller of the good. The social planner’s optimal price schedule differs from the laissez-faire price schedule $P^{\text{LF}}(q) = cq$ whenever $\omega(\theta) \not\equiv \alpha$.*

In other words, whenever the social planner’s objective differs from the utilitarian objective, the social planner can strictly improve on the laissez-faire outcome using a mix of taxation and subsidies. We prove Proposition 1 in Appendix C.4.

Opt-In Benchmark. In some redistribution programs (e.g., public housing and some public childcare and education programs), it may be difficult for consumers who receive subsidies to purchase additional units of the good in the private market. To model this situation, suppose that the social planner can produce units of the good and sell them directly to consumers and that she can require consumers to *opt-in* to the subsidy program, in which case the participating consumer cannot purchase additional units of the good in the private market. In that case, the social planner can charge more the private market price for some units of the good, as long as the *average* price of any quantity of the good sold in the subsidy program is lower than the private market price. This expands the set of feasible mechanisms that can be implemented by the social planner.

In a companion paper, we study the resulting subsidy design problem and characterize the scope of market intervention as follows.

Proposition 2 (scope of intervention with opt-in; [Kang and Watt, 2024](#)). *Suppose the social planner can require consumers to participate in either a subsidized market for the good or the private market for the good (but not both). Then the social planner offers subsidized quantities of consumption if and only if $\max_{\theta} \omega(\theta) > \alpha$.*

Intuitively, when the social planner can prohibit consumers from topping up their subsidized allocations in the private market, the social planner can always target higher welfare weight consumers by offering small subsidies near their laissez-faire consumption levels without those subsidies attracting lower welfare weight consumers to opt into the subsidy program. Whereas [Lemma 3](#) implies that consumption subsidies (without restrictions on topping up) must increase in type for all $\theta \in \Theta$, that property need only hold for a strict subset of consumers opting into the subsidy program when the planner can enforce private market restrictions. This expands the scope of redistributive intervention using subsidies.

Linear Subsidies. Suppose that the social planner is constrained to choosing linear subsidies, which are subsidies for which $\sigma(q) = sq$ for all $q \in [0, A]$ and for some $s \in [0, 1]$. We now study the implications of that constraint on the scope of redistribution via subsidies.

For this section only, we make two additional assumptions that simplify the statement of our results:

Assumption 1. *The valuation function is quadratic: $v(q) = Aq - \frac{1}{2}q^2$.*

[Assumption 1](#) implies that each consumer has linear demand curves, $D(p, \theta) = A - \frac{p}{\theta}$.

Assumption 2. *The parameters c and $\underline{\theta}$ are such that $D(c, \underline{\theta}) = 0$.*

[Assumption 2](#) implies that there is a subset of consumers who choose not to purchase any quantity of the good in the laissez-faire economy.

Proposition 3. *Under [Assumption 1](#) and [2](#), a social planner restricted to choosing $\sigma(q) = sq$ for some $s \in [0, 1]$ intervenes by setting $s < 1$ if and only if*

$$\int_{\underline{\theta}}^{\bar{\theta}} \frac{\int_{\underline{\theta}}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)}{\theta^2} d\theta > 0.$$

Compared to Theorem 1, where the social planner intervenes whenever the maximum value of $\int_{\hat{\theta}}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)$ over $\hat{\theta} \in \Theta$ is positive, a social planner restricted to using linear subsidies intervenes only if the $\frac{1}{\theta^2}$ -weighted average of those terms is positive, which is clearly strictly more restrictive. Intuitively, the social planner considers an average of this term—which captures the social planner’s incentive to distort a type’s consumption—across all consumers because linear subsidies always distort the consumption of each consumer type. On the other hand, with nonlinear subsidies, the social planner can design a subsidy program that distorts the consumption only of types for whom that term is strictly positive. As a result, Proposition 3 captures the loss associated with studying linear subsidies alone: there may be many markets for which intervention with nonlinear subsidies is justified where linear subsidies would not be beneficial. The proof of Proposition 3 is in Appendix C.5, along with analysis of optimal linear subsidies in the general case without Assumption 1 and 2, for which the results are qualitatively similar, but the expression for the condition for intervention is more complicated.

4 Optimal Mechanism

In this section, we explain the second main result of this paper, which is an explicit characterization of the optimal subsidy mechanism.

4.1 Preliminaries

Before introducing our main theorem, we discuss some additional notation and preliminary results needed for its statement.

Virtual Welfare. We define the *virtual welfare* $J(\theta)$ of a type θ as follows:

$$J(\theta) = \theta + \frac{\int_{\theta}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)}{\alpha f(\theta)} + \frac{\max\{\mathbf{E}[\omega] - \alpha, 0\} \cdot \underline{\theta} \delta_{\underline{\theta}}(\theta)}{\alpha f(\theta)},$$

where $\delta_{\underline{\theta}}(\theta)$ is a point mass (Dirac delta) centered at $\underline{\theta}$.

As is standard in mechanism design, the virtual welfare J captures the contribution to the social planner’s objective of a dollar-equivalent of consumer surplus allocated to type θ , accounting for the information rents accruing to types above θ as a consequence of the (IC) constraint. Our expression for J also incorporates the (NLS) constraint, allowing us to account for its possible interactions with the lower-bound constraint in our solution.

As a result, the expression for virtual welfare consists of three additive terms, capturing the social planner’s tradeoff between equity and efficiency. The first term represents the social planner’s benefits from efficiently allocating the good. The second term captures the redistributive benefits of allocating a unit of the good to type θ , which depends on the social planner’s weights on consumers of all types above θ , as any increase in utility for type θ must be granted to all higher types.¹⁰ The third term captures the (NLS) constraint, which binds when $\mathbf{E}[\omega] > \alpha$, leading to a point mass at $\underline{\theta}$. The consumer with type $\underline{\theta}$ plays an important role in our results because satisfying (NLS) and individual rationality for type $\underline{\theta}$ implies those constraints are satisfied for all higher types. When $\mathbf{E}[\omega] > \alpha$, the social planner would like to make a cash transfer to all consumers, but cannot by (NLS), so instead offers all consumers a free quantity of the good.

We refer to $J(\theta) - \theta$ as the *distortion term* of the virtual welfare function. A positive distortion term means the social planner values the consumption of type θ more than the consumer: she wants to distort his consumption upwards. Note that $J(\bar{\theta}) = \bar{\theta}$, reflecting the classic “no distortion at the top” property.

Ironing. We define the ironing operator $\bar{\phi} : [\underline{\theta}, \hat{\theta}] \rightarrow \mathbb{R}$ applied to a generalized function ϕ on an arbitrary interval $[\underline{\theta}, \hat{\theta}] \subseteq \mathbb{R}$ (cf. Myerson, 1981; Toikka, 2011) as follows. Let $\Phi : [\underline{\theta}, \hat{\theta}] \rightarrow \mathbb{R}$ be defined by $\Phi(\theta) = \int_{\underline{\theta}}^{\theta} \phi(s) dF(s)$, and let $\text{co } \Phi$ be the *convex envelope* of Φ , which is the pointwise largest convex function satisfying for all $\theta \in [\underline{\theta}, \hat{\theta}]$, $\text{co } \Phi(\theta) \leq \Phi(\theta)$. Then $\bar{\phi}$ is the monotone function satisfying

$$\text{for all } \theta \in [\underline{\theta}, \hat{\theta}], \quad \int_{\underline{\theta}}^{\theta} \bar{\phi}(s) dF(s) = \text{co } \Phi(\theta).$$

We illustrate this ironing operation in Figure 3.

Several properties of the ironing operator are used in our results.

First, for any $s \in [\underline{\theta}, \theta]$, either $\bar{\phi}(s) = \phi(s)$ or s is contained in an *ironing interval*, on which $\bar{\phi}$ is constant and equal to that interval’s F -weighted average of ϕ . This result is well-known but established formally in Lemma 5 in Appendix A.

Second, note that the definition of $\bar{\phi}$ depends on the domain of ϕ through the construction of $\text{co } \Phi$, as illustrated in Figure 3. Our results apply the ironing operator to the restriction of the virtual welfare to subintervals $[\underline{\theta}, \hat{\theta}] \subseteq \Theta$, denoted $J|_{[\underline{\theta}, \hat{\theta}]}$. One important property of restricted-domain ironing that we establish in Appendix A is that the value of an ironed function at θ

¹⁰ Note that the second term is zero for a utilitarian social planner (with $\omega(\theta) \equiv \alpha$). By setting $\omega(\theta) \equiv 0$ and $\alpha = 1$, we obtain the profit-maximizing social planner, and $J(\theta)$ is the standard virtual value.




Figure 3: Illustrating the ironing operation applied to ϕ on domains $[\underline{\theta}, \hat{\theta}]$ (in red) and $[\underline{\theta}, \hat{\theta}']$ (in black). The ironed functions agree on $[\underline{\theta}, \theta_1]$, after which $\overline{\phi}|_{[\underline{\theta}, \hat{\theta}]} > \overline{\phi}|_{[\underline{\theta}, \hat{\theta}]}$.

can decrease but not increase as the domain extends rightward, with strict decreases occurring if and only if θ lies in an ironing interval of the larger domain. That is, for $\hat{\theta}' > \hat{\theta}$, we have $\overline{\phi}|_{[\underline{\theta}, \hat{\theta}]}(\theta) \geq \overline{\phi}|_{[\underline{\theta}, \hat{\theta}']}(\theta)$, as illustrated in Figure 3.

In mechanism design, ironing is used to address violations of the monotonicity constraint. In the subsidy design problem we study, the monotonicity constraint interacts with the lower-bound constraint, which requires us to identify a new ironing procedure accounting for these interactions. While ironing is typically needed in mechanism design because of “irregularities” in the type distribution, that is not the only possible source of ironing in the subsidy design problems. The virtual welfare function can also have nonmonotonicities caused by the point mass at $\underline{\theta}$ that arises whenever $\mathbf{E}[\omega] > \alpha$ (as discussed above) or nonmonotonicities in $\omega(\theta)$.

4.2 Characterization of the Optimal Mechanism

With these preliminaries, we now state our second main result.

Theorem 2 (characterization of the optimal allocation). *The optimal subsidy allocation rule is unique, continuous in θ , and satisfies*

$$q^*(\theta) = D(c, H(\theta)),$$

where $H(\theta)$ is the subsidy type of type θ , defined as

$$H(\theta) = \begin{cases} \theta & \text{if } \overline{J}_{|\underline{\theta}, \theta]}(\theta) \leq \theta, \\ \overline{J}_{|\underline{\theta}, \kappa_+(\theta)]}(\theta) & \text{otherwise,} \end{cases}$$

and $\kappa_+(\theta) = \inf \left\{ \hat{\theta} \in \Theta : \hat{\theta} \geq \theta, \text{ and } \hat{\theta} \geq \overline{J}_{|\underline{\theta}, \hat{\theta}]}(\hat{\theta}) \right\}$ or $\bar{\theta}$ if that set is empty.


We derive Theorem 2 in Appendix A, along with an extension that permits more general lower-bound constraints.

In the optimal subsidy mechanism, the social planner induces a consumer of type θ to demand the same quantity as a consumer of type $H(\theta)$ in the laissez-faire mechanism. We call $H(\theta)$ the *subsidy type* of type θ . Given the expression $q^*(\theta) = D(c, H(\theta))$ and the strict monotonicity of demand in type, there is a one-to-one correspondence between subsidy type and quantity, so it is equivalent for the social planner to choose $H(\theta)$ rather than $q(\theta)$. Theorem 2 clarifies that this transformed type space is the right space in which to conduct analysis of this problem.


Treating $H(\theta)$ as the decision variable, the constraints on q can be transformed to constraints on H . First, the requirement that q be nondecreasing in θ (from (IC)) is equivalent to H being nondecreasing in θ by the strict monotonicity of demand in type. Second, the requirement that $q(\theta) \geq q^{\text{LF}}(\theta)$ is equivalent to $H(\theta) \geq \theta$, again by the monotonicity of demand in type. The continuity of q^* (which is not a constraint but a property of q^* claimed in Theorem 2) is equivalent to the continuity of H . An important part of the proof of Theorem 2 is establishing those properties for the function H described in the statement of the theorem, with the proof exploiting monotonicity properties of the domain-restricted ironing operator.

Constructing the Subsidy Type. Figure 4 illustrates the construction of the subsidy type $H(\theta)$ for two possible virtual welfare functions. Constructing H requires a “double ironing” calculation: one ironing operation is used to determine where the lower bound is binding (based on $J_{|\underline{\theta}, \theta]}$) and then another to determine the optimal value of H given the ironing intervals determined in the first step (based on $[\underline{\theta}, \kappa_+(\theta))$). In each panel, we plot the virtual welfare $J(\theta)$ (the red dashed curve) and the forty-five degree line (the blue dashed line), which is a lower bound on the choice of $H(\theta)$.

In the example illustrated in Figure 4(a), the virtual welfare $J(\theta)$ is increasing in θ , so $\overline{J}_{|\underline{\theta}, \theta]}(\theta) = J(\theta)$ for each θ . To construct the subsidy type, Theorem 2 implies that we divide the type space into intervals according to whether $\overline{J}_{|\underline{\theta}, \theta]}(\theta) = J(\theta)$ is larger or smaller than θ . In this



(a) In this case, $J(\theta)$ is increasing, and the subsidy type $H(\theta)$ is equal to the lower bound θ for $\theta \leq \theta_1$ and then $J(\theta)$ for $\theta \geq \theta_1$.



(b) In this case, $J(\theta)$ is nonmonotone with a point mass at $\underline{\theta}$, and the subsidy type $H(\theta)$ is equal to $\overline{J|_{[\underline{\theta}, \hat{\theta}]}}$ for $\theta \leq \hat{\theta}$ and then θ for $\theta \geq \hat{\theta}$.

Figure 4: Constructing the subsidy type $H(\theta)$

case, $\overline{J|_{[\underline{\theta}, \theta]}}(\theta) < \theta$ for $\theta \leq \theta_1$, so the subsidy type on $[\underline{\theta}, \theta_1]$ equals θ . On the other hand, for $\theta \geq \theta_1$, Theorem 2 implies that the subsidy type is $\overline{J|_{[\underline{\theta}, \theta]}}(\theta) = J(\theta)$. This implies that the optimal allocation rule in this case is $q^*(\theta) = q^{\text{LF}}(\theta)$ for $\theta \leq \theta_1$ and $q^*(\theta) = D(J(\theta))$ for $\theta \geq \theta_1$.

In the example illustrated in Figure 4(b), the virtual welfare is nonmonotone with a point mass at $\underline{\theta}$, depicted as a red dot in the diagram and corresponding, as described above, to the case in which $\mathbf{E}[\omega] > \alpha$. In this case, $\overline{J|_{[\underline{\theta}, \theta]}}(\theta)$ exceeds θ and is decreasing in θ over $[\underline{\theta}, \hat{\theta}]$ until it hits the lower bound at $\hat{\theta}$. This means that for all $\theta \leq \hat{\theta}$, the subsidy type is equal to J ironed over the interval $[\underline{\theta}, \hat{\theta}]$, which is simply a constant equal to the F -weighted average of J on $[\underline{\theta}, \hat{\theta}]$. On the other hand, for all $\theta \geq \hat{\theta}$, we have $\overline{J|_{[\underline{\theta}, \theta]}}(\theta) < \theta$, so $H(\theta) = \theta$. As a result, $q^*(\theta) = q^{\text{LF}}(\hat{\theta})$ for all $\theta \leq \hat{\theta}$, and $q^*(\theta) = q^{\text{LF}}(\theta)$ for all $\theta \geq \hat{\theta}$.

Role of the Lower Bound Constraint. To understand the role of consumers' access to the private market in the subsidy design problem, we now compare the allocation rule q^* to the optimal allocation rule when the social planner can shut down the private market and is thus not subject to the lower bound constraint (LB). This is equivalent to asking about the effect of the social planner's inability to impose nonlinear taxes on consumption in our model. Broadly, we find that the social planner's inability to levy nonlinear taxes causes the social planner to levy more generous subsidies. That is, optimal subsidies are not merely found by identifying the optimal tax-and-subsidy schedule and setting all taxes to zero.

To study this question, we compare the optimal allocation rule in the subsidy design problem

to the optimal allocation rule in the shutdown problem (derived in Appendix C.4 as (cf. Toikka, 2011))

$$q^{\text{SD}}(\theta) = D(c, \bar{J}(\theta))$$

Note that as long as J is not almost everywhere the identity function (which corresponds to the case in which $\omega(\theta) = \alpha$ almost everywhere), q^{SD} differs from q^{LF} . In that case, a social planner who can shut down the private market (or levy taxes) finds it beneficial to intervene in the market, reflecting our finding, discussed in Section 3.2, that the consumers' ability to access the private market restricts the scope of redistribution by the social planner.

As is clear from Figure 4, the subsidy type in our problem is *not* obtained simply as the pointwise maximum of $\bar{J}(\theta)$ (the gray curve) and the lower bound θ (the blue dashed line). The basic reason for this is that the monotonicity constraint and the (LB) constraint may interact: although the two constraints cannot both bind on the same interval (because q^{LF} is monotone), the two constraints may interact in determining *where* each constraint binds.

Comparing the shutdown benchmark allocation q^{SD} to the subsidy solution q^* described in Theorem 2, we see that private market access distorts consumption upwards, with q^* strictly exceeding q^{SD} wherever $H(\theta) > \bar{J}(\theta)$. That distortion affects a larger set of agents than those for whom the lower-bound constraint would be binding at q^{SD} . In Figure 4(b), for example, $q^* > q^{\text{SD}}$ for all consumer types, even though the (LB) constraint would be binding on q^{SD} only on $[\theta_1, \bar{\theta}]$. In other words, the optimal allocation is not obtained by relaxing (LB) and then enforcing it on the solution to the relaxed problem.

The intuition for this upward distortion can be understood in terms of information rents: the cost to a consumer of type θ associated with reporting a lower type θ' is reduced in the presence of the private market because he can always purchase additional units of the good in the private market, resulting in higher information rents accruing to type θ . That effect is present for all subsidized types, not only those for whom the private market constraint would be binding in q^{SD} .

Mathematically, taking the pointwise maximum of q^{SD} and q^{LF} is not optimal if the (LB) constraint binds on a subset of an ironing interval of q^{SD} , because that would violate a pooling condition for optimality. In Figure 4, for example, if $q(\theta) = \max\{q^{\text{LF}}(\theta), q^{\text{SD}}(\theta)\}$, consumers in $[\underline{\theta}, \theta_1]$ would be allocated less than is optimal for the social planner because the average virtual welfare of those types is higher than $\bar{J}(\theta)$ on that interval.

4.3 Discussion of Proof Approach

There are three main steps in our proof of Theorem 2.

The first step is to rewrite the subsidy design problem as a convex program with a lower-bound constraint, including deriving the expression for $J(\theta)$ above. This step applies standard mechanism design techniques to reformulate the constraints and objective in terms of the allocation rule. We deviate from standard mechanism design approaches by incorporating the (NLS) constraint into the virtual welfare, allowing us to account for interactions between it and the other constraints.

The second step involves establishing the feasibility of q^* . This step exploits several properties of the ironing operator, which extend similar properties established for a discrete ironing operator used in the statistics literature on isotonic regression (cf. [Van Eeden, 1956](#); [Barlow, Bartholomew, Bremner and Brunk, 1972](#); [Robertson, Wright and Dykstra, 1988](#)). These propositions, derived in [Appendix A](#), may be of independent interest in the generalized ironing literature.

The third step is to verify the optimality of q^* . We begin with a standard approach in mechanism design (see, e.g., [Toikka, 2011](#)), relating the subsidy design problem’s solution to that of a simpler convex program that can be solved pointwise. The challenge of that approach is to determine the related problem and to verify that the two problems share an optimizer. As discussed further in [Kang \(2024\)](#), this step may be thought of as a projection of the virtual welfare J onto the space of feasible functions, which is complicated here by the fact that the feasible set is not a convex cone (as it is in the absence of the lower-bound constraint). While [Yang and Zentefis \(2024\)](#) characterize the extreme points of the feasible set (which they call a “monotone function interval”), the optimizer of our convex program is typically an interior point of that set. Whereas [Corrao et al. \(2023\)](#) derive certain properties of the optimizer of a similar convex program arising in a different setting, our techniques allow us to obtain an explicit characterization of the optimizer for separable consumer preferences à la [Mussa and Rosen \(1978\)](#). To verify the solution we identify, we exploit a variational inequality characterizing the condition for optimality in (OPT), which differs from the more standard Lagrangian approach used to solve mechanism design programs in the literature (for example, by [Amador et al., 2006](#) and [Amador and Bagwell, 2013](#)).

5 Implications for Subsidy Design

In this section, we discuss the implications of our main results for subsidy design, providing a detailed characterization of the optimal subsidy mechanism under additional assumptions on the market’s primitives. In particular, we discuss how the main features of the optimal subsidy mechanism depends on two key factors: whether the opportunity cost of funds is high or low, and

whether welfare weights are positively or negatively correlated with type.¹¹ Table 1 summarizes the optimal subsidies in each case.

	High cost of funds ($\mathbf{E}[\omega(\theta)] \leq \alpha$)	Low cost of funds ($\mathbf{E}[\omega(\theta)] > \alpha$)
Negative Correlation ($\omega(\theta)$ decreasing)	$q^*(\theta) = q^{\text{LF}}(\theta)$ for all $\theta \in \Theta$.	For some $\theta_\alpha \in \Theta$: $q^*(\theta) \geq q^{\text{SD}}(\theta)$ for all $\theta \leq \theta_\alpha$, $q^*(\theta) = q^{\text{LF}}(\theta)$ for all $\theta \geq \theta_\alpha$.
Positive Correlation ($\omega(\theta)$ increasing)	For some $\theta_\alpha \in \Theta$: $q^*(\theta) = \begin{cases} q^{\text{LF}}(\theta) & \text{for } \theta \leq \theta_\alpha, \\ q^{\text{SD}}(\theta) & \text{for } \theta \geq \theta_\alpha. \end{cases}$	$q^*(\theta) = q^{\text{SD}}(\theta)$ for all $\theta \in \Theta$.

Table 1: Dependence of the optimal mechanism on market primitives

The first factor—how the opportunity cost of funds compares to the average welfare weight—determines whether the optimal subsidy mechanism involves *public provision*, by which we mean a free endowment of the good for all consumers, as we show in Section 5.1 below. The second factor—whether welfare weights are positively or negatively correlated with type—affects the subsidy design by changing the sign of the distortions in the virtual welfare term. In particular, when ω is increasing in θ , the distortion term of the virtual welfare changes sign at most once from negative to positive, with the initial sign of the distortion term pinned down by the sign of $\mathbf{E}[\omega] - \alpha$. When ω is decreasing in θ , the distortion term changes sign at most once from positive to negative, and so the social planner wants to distort the consumption of lower types upwards. We show how the social planner can most effectively use in-kind subsidies to target the consumers with positive distortion terms in each case.

5.1 When Should Non-Market Allocations Be Used?

In some subsidy markets, governments provide all eligible consumers a baseline quantity of the good. The following proposition sets forth when such non-market allocations are optimal in our model.

¹¹ While our main result, Theorem 2, does not require additional assumptions on ω , the sharp characterization of the optimal subsidy mechanism we derive in this section exploit monotonicity. In Appendix D.2, we extend this detailed analysis to the case in which welfare weights are *U-shaped* or *inverted U-shaped* as a function of θ .

Proposition 4 (non-market allocations). *Non-market allocations arise in the optimal subsidy mechanism depending on the sign of $\mathbf{E}[\omega] - \alpha$, as follows:*

- (a) *If $\mathbf{E}[\omega] > \alpha$, the social planner provides all consumers $q^*(\underline{\theta})$ units of the good for free.*
- (b) *If $\mathbf{E}[\omega] < \alpha$, there is no public provision of the good to all consumers, and the initial $q^*(\underline{\theta})$ are sold at the laissez-faire price, c .*
- (c) *If $\mathbf{E}[\omega] = \alpha$, the social planner is indifferent between providing $q^*(\underline{\theta})$ units of the good for free and charging any price less than or equal to $cq^*(\underline{\theta})$.*

The proof of Proposition 4 is in Appendix A.

Intuitively, when $\mathbf{E}[\omega] > \alpha$, the social planner would like to make a cash transfer to all consumers but is constrained by the (NLS) constraint. In that case, she never charges a positive price for the initial units of the good, since making that quantity free is equivalent to providing a cash transfer to all consumers. On the other hand, when $\mathbf{E}[\omega] < \alpha$, the social planner would prefer to tax consumers but cannot by assumption. In that case, by Theorem 2, $q^*(\underline{\theta}) = q^{\text{LF}}(\underline{\theta})$, and if the planner ever offered the consumer with type $\underline{\theta}$ a subsidy, an increase in $t^*(\underline{\theta})$ to the laissez-faire payment level $cq^{\text{LF}}(\underline{\theta})$ would function like a tax on all consumers and benefit the social planner.

To determine payments for units beyond the first $q^*(\underline{\theta})$, a payment schedule implementing the optimal subsidy mechanism can be determined using the taxation principle (cf. Hammond, 1979; Guesnerie, 1981). We provide an explicit expression of the optimal payment schedule in Proposition 15 in Appendix A and offer an intuitive description here. Given the payments for the lowest type determined in Proposition 4, the allocation rule $q(\cdot)$ determines the slope of the payment schedule $P^\sigma(z)$ via the envelope theorem (cf. Milgrom and Segal, 2002). In order to induce a consumer of type θ to consume a quantity $q(\theta)$, the slope of $P^\sigma(z)$ at $z = q(\theta)$ must equal the marginal utility of the consumer of type θ , equal to $\theta v'(q(\theta))$. Because $q(\theta) \geq q^{\text{LF}}(\theta)$ and $q^{\text{LF}}(\theta)$ is chosen so the marginal utility of consumption is c , the slope of $P^\sigma(z)$ is always bounded above by c (confirming that total subsidies are increasing in quantity) and strictly so when $H(\theta) > \theta$.

5.2 How Do Optimal Subsidies Depend on Quantity Consumed?

We now describe the structure of the optimal subsidy mechanism in two benchmark cases: when welfare weights are positively correlated with demand and when they are negatively correlated.

Negative Correlation. Suppose that demand and welfare weights are negatively correlated, in that the social planner's welfare weights are a decreasing function of the consumers' types. The distortion term of the virtual welfare at θ , which determines the social planner's incentive to distort θ 's consumption upwards or downwards, depends on the expected welfare weight of all consumers with higher demand than θ . In this case, if $\mathbf{E}[\omega(\theta)|\theta \geq \hat{\theta}] \leq \alpha$ for any type $\hat{\theta}$, then the decreasing property of ω implies that the distortion term is negative for all consumers with types higher than $\hat{\theta}$ as well. As a result, the distortion term of the virtual welfare changes sign at most once from positive to negative (as in Figure 4(b), for example), with the initial sign of the distortion pinned down by the sign of $\mathbf{E}[\omega] - \alpha$. This leads to the following structure of the optimal subsidy mechanism.

Proposition 5 (optimal mechanism with negative correlation). *Suppose ω is decreasing in θ . Then:*

- (a) *If the opportunity cost of funds is high, so $\mathbf{E}[\omega] \leq \alpha$, the social planner offers no subsidies, and $q^*(\theta) = q^{\text{LF}}(\theta)$.*
- (b) *If the opportunity cost of funds is low, so $\mathbf{E}[\omega] > \alpha$, there exists a type $\theta_\alpha \in \Theta$ such that*

$$q^*(\theta) = \begin{cases} D(c, \overline{J|_{[\underline{\theta}, \theta_\alpha]}}(\theta)) & \text{if } \theta \leq \theta_\alpha \\ q^{\text{LF}}(\theta) & \text{if } \theta > \theta_\alpha. \end{cases}$$

If $\min \omega \geq \alpha$, then $\theta_\alpha = \bar{\theta}$, otherwise $\theta_\alpha < \bar{\theta}$. This allocation is implemented using a free endowment of the good to all consumer types and possibly additional discounts for consumption up to a maximum level $q^(\theta_\alpha)$.*

The proof of Proposition 5 is in Appendix B.

Consistent with Theorem 1, the social planner does not intervene when $\mathbf{E}[\omega] \leq \alpha$. The social planner offers subsidies only if the opportunity cost of funds is low, which is exactly the case in which the no lump-sum transfer constraint is binding. In that case, the social planner would prefer to make a cash transfer to all consumers. The marginal prices as a function of quantity in each case are illustrated in Figure 5.

Positive Correlation. Suppose that demand and welfare weights are positively correlated, in that the social planner's welfare weights are an increasing function of the consumers' types. In that case, if for any type $\hat{\theta}$, we have $\mathbf{E}[\omega(\theta)|\theta \geq \hat{\theta}] \geq \alpha$, then the distortion term is positive for all




Figure 5: Marginal prices as a function of quantity for decreasing welfare weights.

consumers with types higher than $\hat{\theta}$. This means that the distortion term of the virtual welfare changes sign at most once from negative to positive (as in Figure 4(a), for example), with the initial sign of the distortion pinned down by the sign of $\mathbf{E}[\omega] - \alpha$. This leads to the following structure of the optimal subsidy mechanism.

Proposition 6 (optimal mechanism with positive correlation). *Suppose ω is increasing in θ . Then:*

(a) *If the opportunity cost of funds is high, so $\mathbf{E}[\omega] \leq \alpha$, then*

$$q^*(\theta) = \begin{cases} q^{\text{LF}}(\theta) & \text{for } \theta \leq \theta_\alpha \\ q^{\text{SD}}(\theta) & \text{for } \theta > \theta_\alpha, \end{cases}$$

where $\theta_\alpha = \bar{\theta}$ if $\max \omega \leq \alpha$, or otherwise θ_α satisfies $\mathbf{E}[\omega | \theta \geq \theta_\alpha] = \alpha$. The optimal mechanism is implemented by increasing subsidies for consumption levels $z > q^{\text{LF}}(\theta_\alpha)$ when $\theta_\alpha < \bar{\theta}$.

(b) *If the opportunity cost of funds is low, so $\mathbf{E}[\omega] > \alpha$, $q^*(\theta) = q^{\text{SD}}(\theta)$, implemented by subsidies for all levels of consumption.*

The proof of Proposition 6 is in Appendix B.

In this case, the social planner intervenes whenever $\max \omega > \alpha$, offering subsidies for consumption beyond a certain minimum level of consumption (possibly the lowest level of

consumption, $q^*(\theta)$, namely when $\mathbf{E}[\omega] > \alpha$). When $\mathbf{E}[\omega] < \alpha$ and $\max \omega > \alpha$, the (NLS) constraint is slack, and the social planner offers in-kind subsidies when she would not otherwise offer cash transfers. This is because only consumers with types higher than the θ_α defined in Proposition 6 receive in-kind subsidies, and the social planner places a higher average welfare weight on those types than the cost of public funds. The marginal prices as a function of quantity are illustrated in each case in Figure 6.




Figure 6: Marginal prices as a function of quantity when demand is positively correlated with need.

5.3 How Do Optimal Nonlinear Subsidies Compare to Linear Subsidies?

As we observed in Section 3.2, the social planner intervenes less often when restricted to linear subsidies compared to when she has access to nonlinear subsidy instruments. In this subsection, we show that the structure of the optimal nonlinear subsidy mechanism differs markedly from linear subsidies.

Our first result clarifies that the optimal nonlinear subsidy mechanism is *never* linear.

Proposition 7. *The optimal subsidy schedule identified in Theorem 2 is never linear: that is, there exists no $s \in [0, 1)$ such that $\sigma^*(q) = cs$.*

This result follows almost immediately as a consequence of the “no distortion at the top” property of the virtual type: any linear subsidy mechanism results in distorted consumption for all consumers, including the one with the highest type, and thus cannot be optimal. The proof of Proposition 7 is in Appendix C.5.

In practice, it may be challenging for the social planner to implement the optimal nonlinear subsidy in some markets, which—as we saw in Theorem 2—may involve a combination of public provision, partial subsidies, and subsidy caps. However, focusing on the negative and positive correlation cases we studied above, we show that it is possible to identify relatively simple improvements over linear subsidies.

Proposition 8. *Suppose the social planner has implemented a subsidy program with $\sigma(q) = cs$ for some $s \in [0, 1)$.*

(a) *Negative correlation: If ω is decreasing in θ , the social planner can always improve over σ by implementing at least one of:*

- (a) *a cap on the subsidy paid to any consumer, so $\sigma'(q) = cs\mathbf{1}_{q \leq \hat{q}}$ for some $\hat{q} \in [0, A]$, or*
- (b) *a free endowment of the good for all consumers, so $\sigma'(q) = c$ for $q \leq \hat{q}$ and $\sigma'(q) = cs$ for $q > \hat{q}$ for some $\hat{q} \in [0, A]$.*

(b) *Positive correlation: If ω is increasing in θ , the social planner can always improve over σ by implementing at least one of:*

- (i) *a floor (or copay) on the subsidy, so $\sigma'(q) = cs\mathbf{1}_{q \geq \hat{q}}$ for some $\hat{q} \in [0, A]$, or*
- (ii) *a free endowment of the good for all consumers, so $\sigma'(q) = c$ for $q \leq \hat{q}$ and $\sigma'(q) = cs$ for $q > \hat{q}$ for some $\hat{q} \in [0, A]$.*

Intuitively, a cap on subsidy payments is beneficial whenever the social planner assigns a lower weight than the cost of public funds to the highest-demand consumers, while a floor on subsidy payments is beneficial if the social planner assigns a lower welfare weight than the cost of public funds for the lowest-demand consumers. The social planner benefits from offering a free quantity of the good whenever the opportunity cost of public funds is lower than the average welfare weight. The proof of Proposition 8 is presented in Appendix C.5 and involves showing that at least one of these improvements must be possible in any market.

5.4 When Does The Planner Want To Restrict Private Market Access?

As we discussed in Section 3.2, it may be possible for the social planner to restrict recipients in some subsidy programs from purchasing additional units in the private market. The question of whether to restrict subsidy recipients from participating in private markets has arisen in many real-world subsidy programs. For example, driven by worries about disparities in educational access

linked to wealth and socioeconomic privilege, China introduced stringent regulations that made it difficult for private tutors to continue operating (Palmer, 2021). As a result, consumers are no longer able to top up their educational consumption through private tutoring and are left to choose between public or private school options. As another example, the German health insurance system requires citizens to choose private or public health insurance coverage, and consumers opting into the private program find it difficult to switch back to the public system (Schmidt-Kasperek, 2023). Similarly, recipients of certain health insurance subsidies called Cost-Sharing Reductions under the Affordable Care Act are limited to consumers selecting “silver” health insurance plans and become ineligible for the subsidies if they opt in to a higher-coverage “gold” plan (Healthcare.gov, 2024).

Restricting consumers’ access to the private market reduces the set of constraints facing the social planner to (IC), (NLS) and an individual rationality constraint, requiring that $U(\theta) \geq U^{\text{LF}}(\theta)$ for each $\theta \in \Theta$. In our model with the (LB) constraint, Lemma 2 clarifies that we need only enforce the individual rationality constraint for $\underline{\theta}$ to ensure that it is satisfied for all types (as we show in Appendix A). That simplification does not apply in the setting where consumers are restricted to opt-in to subsidy programs (and opt-out of private markets). Instead, we show in our companion paper (Kang and Watt, 2024) that the individual rationality constraint imposes a majorization constraint on the allocation rule, requiring that

$$\underline{U} + \int_{\underline{\theta}}^{\theta} v(q(s)) \, ds \geq \underline{U}^{\text{LF}} + \int_{\underline{\theta}}^{\theta} v(q(s)) \, ds \quad \text{for all } \theta \in [\underline{\theta}, \bar{\theta}].$$

In our companion paper, we show that—in the case ω and θ are positively correlated—the optimal subsidy mechanism in this case satisfies the (LB) constraint. As a result, there is no benefit to restricting private market access in the case of positive correlation. This means that a positive correlation between demand and welfare weights not only ensures that in-kind subsidies are self-targeting, but also reduces the need to enforce private market restrictions on consumers.

On the other hand, when ω and θ are negatively correlated, the optimal subsidy mechanism in the opt-in model improves over the laissez-faire whenever $\max \omega > \alpha$. As a result, whenever $\max \omega > \alpha > \mathbf{E}[\omega]$, the social planner strictly prefers the mechanism with opt-in to the mechanism identified in this paper. Intuitively, if the social planner can prohibit subsidy recipients from purchasing additional units of the good in the private market, she can offer a subsidy tied to low levels of consumption, and, as long as that subsidy is not too generous, high-demand consumers would be unwilling to report a low type to receive the subsidy. This means that the social planner

can target high-need consumers without also subsidizing consumers with lower need (but higher demand for the good).

5.5 How Do Optimal Subsidies Depend on Market Primitives?

We now study the comparative statics of the optimal subsidy program in the economic primitives of the market. In particular, we will study the effect of three changes on the optimal subsidy program: (a) a pointwise increase in ω or a decrease in α , corresponding to a secular increase in the social planner's desire to redistribute to each consumer, (b) a decrease in ω in the majorization order, corresponding to increased positive correlation between ω and θ , and (c) a change in the marginal cost of production c or the consumer's preferences over consumption v .

While we have framed these questions in terms of changes in economic primitives, the results in this section also shed light on subsidy choices not explicitly included in our model, including product choice, eligibility restrictions and costly screening. We briefly touch on these questions in this section, but explore this perspective further in Section 7.2 below.

Change in Preferences for Redistribution to All Types. We first study the effect of an increase in the social planner's preference to redistribute to all consumer types, as may be caused by a change in government.

Proposition 9 (increasing preferences for redistribution). *Suppose that there is a shift in the preference for redistribution from ω to $\tilde{\omega}$ with $\tilde{\omega}(\theta) \geq \omega(\theta)$ for all $\theta \in \Theta$, and/or a reduction in the opportunity cost of funds from α to $\tilde{\alpha} < \alpha$. Then, the optimal subsidy mechanism leads to a higher total weighted surplus and is more generous to recipients, in that:*

- (a) each consumer's subsidy type increases,¹² to $\tilde{H}(\theta) \geq H(\theta)$,
- (b) the set of subsidized consumers increases (in the sense of set inclusion),
- (c) each consumer's total allocation increases, so $\tilde{q}^*(\theta) \geq q^*(\theta)$,
- (d) the total subsidy received by each consumer increases, and
- (e) each consumer is better off overall, so $\tilde{U}^*(\theta) \geq U^*(\theta)$.

¹² In this section, we use the word "increase" to refer to weak improvements: that is, when we say that a variable x "increases," its new value \tilde{x} satisfies $\tilde{x} \geq x$.

We prove Proposition 9 in Appendix C.6.

Intuitively, an increased preference for redistribution causes the virtual welfare of each type to increase to $\tilde{J}(\theta) > J(\theta)$, resulting in a stronger preference for the social planner to distort consumption all types upwards. Proposition 9 establishes that this change leads the social planner to increase the subsidy and consumption distortion of existing subsidy recipients (the intensive margin) and to subsidize types that did not previously receive subsidies (the extensive margin), as illustrated in Figure 7 (separately for the case in which ω is decreasing in θ and when it is increasing).

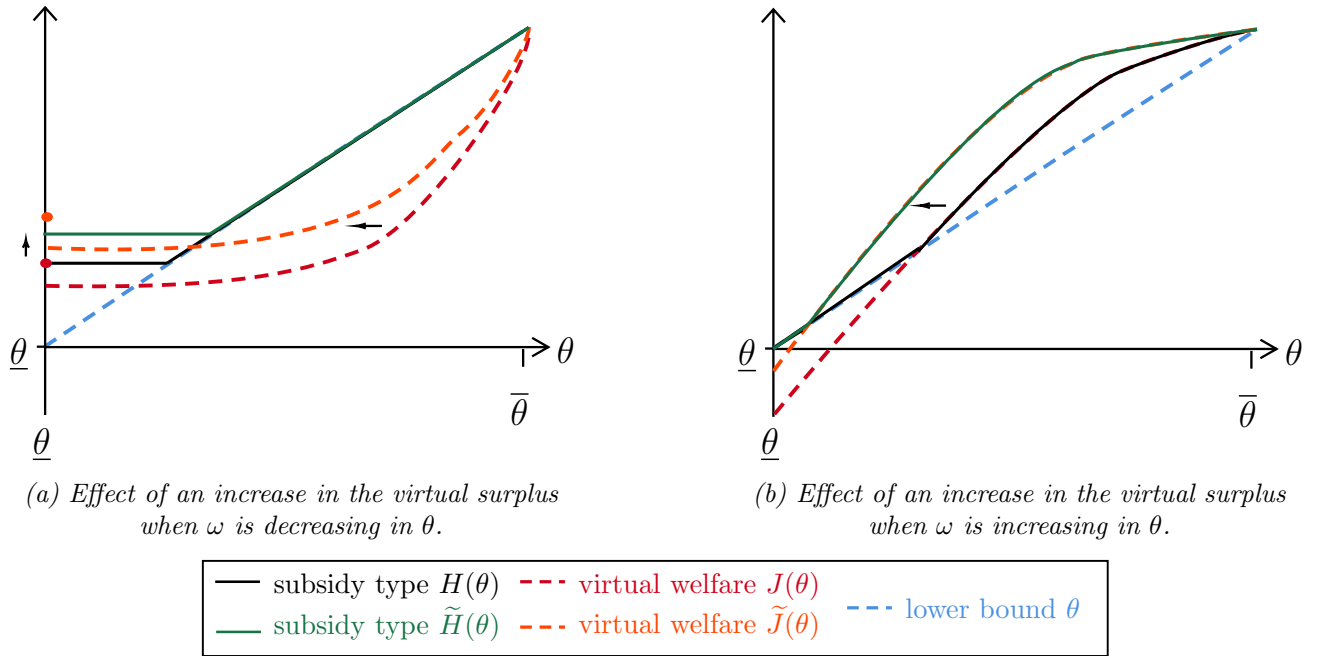


Figure 7: Effects of a change in the virtual surplus on the subsidy type

This comparative static also sheds light on differences in the optimal subsidy mechanism for different groups of consumers with the same demand for the good. Holding all else equal, Proposition 9 suggests that the social planner and eligible consumers both prefer subsidies offered to groups of consumers with higher welfare weights as a function of demand.

Changing Correlation We now consider the effect on the optimal subsidy mechanism of an increase in the correlation between ω and θ , leaving $\mathbf{E}[\omega]$ unchanged. Such a change in preferences might also be caused a change in government, with the social planner assigning stronger welfare weights to lower- or higher-demand consumers.

Formally, we study the effect of a *decrease* in $\omega(\cdot)$ in the *majorization order* (cf. [Hardy, Littlewood and Pólya, 1934](#); [Kleiner, Moldovanu and Strack, 2021](#)). Welfare weights $\tilde{\omega}$ are majorized by ω if for all $\hat{\theta} \in \Theta$, we have

$$\mathbf{E}[\tilde{\omega}(\theta) \mid \theta \geq \hat{\theta}] \geq \mathbf{E}[\omega(\theta) \mid \theta \geq \hat{\theta}],$$

and $\mathbf{E}[\tilde{\omega}] = \mathbf{E}[\omega]$. In other words, the social planner places more weight on high-demand consumers under welfare weights $\tilde{\omega}$ than ω . A decrease in the majorization order leads to a pointwise increase in the average welfare weight above any type, implying a pointwise increase in the distortion term and thus virtual welfare, as illustrated in Figure 8.




Figure 8: Increasing interdependence: $\tilde{\omega}$ majorized ω , leading to a decrease in the distortion term.

A similar analysis as for Proposition 9 implies the following effect of increasing correlation between demand and need on the optimal subsidy mechanism.

Proposition 10 (effect of changing interdependence). *Suppose that $\tilde{\omega}$ is majorized by ω . Then, the total weighted surplus is greater and the optimal subsidy mechanism is more generous given $\tilde{\omega}$ than ω , in that:*

- (a) each consumer's subsidy type increases, to $\tilde{H}(\theta) \geq H(\theta)$,
- (b) the set of subsidized consumers is larger (in the sense of set inclusion),
- (c) each consumer's total allocation increases, so $\tilde{q}^*(\theta) \leq q^*(\theta)$,
- (d) the total subsidy received by each consumer increases, and
- (e) each consumer's utility increases, so $\tilde{U}^*(\theta) \leq U^*(\theta)$.

We prove Proposition 10 in Appendix C.6.

This comparative static also allows us to compare the optimal subsidy mechanisms for different products with the same demand system but different relationships between demand and social preferences. Holding all else equal, Proposition 10 suggests that the social planner and the average eligible consumer prefer subsidies for goods for which demand is more positively correlated with welfare weights.

Changes in Demand or Costs. Finally, we study an increase in the laissez-faire demand for the good, as might be caused by a reduction in the marginal cost of the good to $\tilde{c} < c$, say, or by a change in the valuation function v to a new valuation function \tilde{v} with $\tilde{v}'(q) \geq v'(q)$ for each $q \in [0, A]$.

Neither change affects the virtual welfare. As a result, they do not affect the calculation of the subsidy type, which does not depend on the cost or the valuation function. Intuitively, the increased demand or reduced cost of the good does not change the designer’s preferences to redistribute utility between agents, which means that the set of subsidized types is unchanged. However, it does increase the social planner’s net benefit from each unit allocated to a consumer, so the optimal subsidy and allocation rule changes, leading to the following.

Proposition 11 (effect of demand increase). *Suppose that there is an increase in the laissez-faire demand for the good caused by a reduction in the cost c or an increase in the marginal utility of consumption v' . Then, the total weighted surplus of the optimal subsidy mechanism increases, the set of subsidized types is unchanged, and the total allocation and consumer surplus of each consumer increases.*

We prove Proposition 11 in Appendix C.6.

Holding all else equal, Proposition 11 suggests that the social planner prefers to subsidize goods with lower production costs or higher marginal values for consumers.

6 Extensions

In this section, we study several extensions of our baseline model to incorporate additional important considerations for subsidy design in some market.

In the first extension, we study the effect of subsidies on prices in the private market. Imperfectly elastic supply can be readily incorporated into our framework by using the subsidy types we calculated above to determine a “subsidized demand curve” and determining the intersection of supply and demand.

In the second extension, we study the effect of taxation programs on the subsidies chosen by the social planner. The existence of exogenous taxation allows the social planner to implement more allocation rules without distortion relative to the efficient consumption benchmark, which increases the weighted surplus achievable via subsidies. When the planner has some limited influence over tax levels, she trades off distortions caused by taxation against improvements in subsidy targeting.

Finally, we consider extensions of our model that endogenize the welfare weights in the social planner's redistributive objective.

6.1 Equilibrium Effects

Until now, we have assumed that introducing the subsidy has no effect on the price of the good in the private market, even though the total consumption—and therefore total production of the good in equilibrium—is increased by the subsidy program. While perfect elasticity of supply may be an appropriate first-order approximation in some markets, the social planner may need to account for the equilibrium effect of the subsidy program on the price of the good.

As we discussed in Section 5.5, an increase in the private market price has two effects on the subsidy design problem: it raises the cost of each unit of subsidized consumption (which, alone, would reduce the optimal subsidy offered by the social planner), but it also allows the social planner to charge higher prices for some units of consumption, while offering larger subsidies on other units, which may improve redistribution.

To incorporate these effects, we now suppose that there is a strictly convex and differentiable industry cost function $C : [0, A] \rightarrow \mathbb{R}$, where $C(Q)$ is the total cost to all suppliers of producing $Q = \int_{\Theta} q(\theta) dF(\theta)$ units of the good. We maintain the assumption that the industry is competitive, so that this cost function leads to an industry supply curve $\Sigma(p) = \max_{Q \in [0, A]} [pQ - C(Q)]$, and write the inverse supply function as $S^{-1}(Q)$ (which is well-defined as a consequence of the differentiability and strict convexity of C).

A change in the private market price affects each consumer's ability to purchase additional units in the private market for the good. We capture this possibility in revised incentive compatibility constraints and individual rationality constraints. For any mechanism (q, t) , let $p = S^{-1}(\int_{\Theta} q(\theta) dF(\theta))$ be the equilibrium price of the good induced by the mechanism, then the revised private market constraint is

$$\text{for all } \theta \in \Theta, q(\theta) \geq D(p, \theta), \tag{PM'}$$

and the revised individual rationality constraint is

$$\text{for all } \theta \in \Theta, \theta v(q(\theta)) - t(\theta) \geq \max_{\hat{q} \in [0, A]} [\theta v(\hat{q}) - p\hat{q}]. \quad (\text{IR}')$$

Proposition 12 (optimal subsidy with equilibrium effects). *The allocation rule solving*

$$\max_{(q,t)} \int_{\Theta} \left[\omega(\theta) [\theta v(q(\theta)) - t(\theta)] + \alpha t(\theta) \right] dF(\theta) - \alpha C \left(\int_{\Theta} q(\theta) dF(\theta) \right),$$

such that (q, t) satisfies (IC), (IR'), (NLS) and (PM'),

is characterized by

$$q^*(\theta) = D(p, H(\theta)), \text{ and } p = S^{-1} \left(\int_{\Theta} q^*(\theta) dF(\theta) \right),$$

where $H(\theta)$ is defined as in Theorem 2.

The proof of Proposition 12 is in Appendix C.7. Note that the definition of the subsidy type $H(\theta)$ is unchanged compared to its definition in Theorem 2. Proposition 12 implies that the optimal allocation rule can be found by constructing a *subsidized demand* curve

$$D(p) = \int_{\Theta} D(p, H(\theta)) dF(\theta)$$

and identifying its intersection p^* with the industry supply curve, determining the competitive market price, as illustrated in Figure 9. Each type θ is then assigned an allocation equal to the demand of an agent with type $H(\theta)$ at the resulting price p^* . Finally, payments are determined via the taxation principle, with the payment of the lowest type set by either the modified (IR') constraint at the prevailing price (when $\mathbf{E}[\omega] < \alpha$) or the no lump-sum transfers constraint (NLS) (when $\mathbf{E}[\omega] \geq \alpha$).

6.2 Taxation of the Private Market

In this section, we discuss the possibility of taxation in the private market. Returning to the perfectly elastic supply setting, we first suppose that the social planner levies an exogenously chosen tax on consumption in the private market and show that this increases weighted total surplus by slackening the constraints in the subsidy design problem. We then discuss the




Figure 9: Illustrating the Subsidized Demand Curve with Equilibrium Effects

implications of that result for a model in which the social planner can choose the tax level while facing a cost of higher taxation in the private market.

Exogenous Taxation Suppose that an exogenously-set tax τ is applied to consumption in the private market, so consumers face a price $(1 + \tau)c$ for each unit of private market consumption. The exogeneity of the tax may result from the tax level being chosen by a different government agency than the one offering the subsidy. We maintain the assumption that the social planner has access to an equally efficient production technology as the private market,¹³ so the subsidy design problem takes the form

$$\max_{(q,t)} \int_{\Theta} \left[\omega(\theta) [\theta v(q(\theta)) - t(\theta)] + \alpha [t(\theta) - cq(\theta)] \right] dF(\theta),$$

subject to (IC), (NLS), an individual rationality constraint incorporating the tax τ ,

$$\theta v(q(\theta)) - t(\theta) \geq \max_{q \in [0,A]} [\theta v(q) - c(1 + \tau)q], \quad (\text{IR}_t)$$

and a lower-bound constraint incorporating the tax τ

$$q(\theta) \geq D((1 + \tau)c, \theta). \quad (\text{LB}_t)$$

¹³ Otherwise, if the social planner must also pay the tax on the production of subsidized units, the subsidy design problem is the same as the baseline model with the price replaced by $c(1 + \tau)$.

Compared to the baseline subsidy design problem studied in this paper, the social planner faces less restrictive individual rationality and lower-bound constraints, allowing her to implement a larger set of allocation rules. Because the objective is the same, exogenous taxation allows the social planner to offer more effective redistribution using in-kind subsidies. Intuitively, exogenous taxation helps the social planner approach the shutdown benchmark in which it has access to both tax and subsidy instruments.

The optimal subsidy mechanism is computed similarly to the baseline model, as described below.

Proposition 13 (optimal subsidy with exogenous taxation). *In the presence of an exogenous tax τ on private consumption, the optimal subsidy allocation is*

$$q^*(\theta) = D(c, H_\tau(\theta))$$

where $H_\tau(\theta) = \overline{H_\tau^\Gamma}(\theta)$ and

$$H_\tau^\Gamma(\theta) = \begin{cases} \frac{\theta}{1+\tau} & \text{if } \overline{J_{[\underline{\theta}, \theta]}}(\theta) \leq \frac{\theta}{1+\tau}, \\ J(\theta) & \text{otherwise.} \end{cases}$$

We prove Proposition 13 in Appendix C.7.

Costly Taxation Now suppose the taxation level τ may be chosen by the social planner. As we discussed in Section 2.4, giving the social planner unconstrained power to set the tax on private consumption makes it possible for her to implement the shutdown benchmark. However, in practice, implementing a tax on the private market would impose costs on consumers outside the subsidy program, which the social planner may account for in the joint tax and subsidy design problem. Supposing that there is a convex cost $\Gamma(\tau)$ associated with taxation in the private market, the social planner now solves

$$\max_{(q, t, \tau)} \int_{\Theta} \left[\omega(\theta) [\theta v(q(\theta)) - t(\theta)] + \alpha [t(\theta) - cq(\theta)] \right] dF(\theta) - \Gamma(\tau),$$

subject to (IC), (NLS), (IR_t) and (LB_t). In that case, the optimal subsidy allocation is determined as in Proposition 13, where the tax level τ^* is chosen to equate the marginal benefits of improved subsidy targeting in the subsidized market against the marginal cost of taxation in the private market. As a result, many of the main features of the optimal subsidy mechanism derived in

Section 4 continue to apply to the setting with costly taxation of the private market.

6.3 Budget Constraints and Endogenous Welfare Weights

In the baseline model, we have taken as given the weights ω the social planner assigns to consumer surplus, as well as the weight α assigned to the cost of subsidy spending. We now discuss possible microfoundations for each of these weights, allowing them to be determined endogenously in the subsidy design problem while maintaining the qualitative features of the optimal mechanisms.

The opportunity cost of subsidy spending, α , can be interpreted as the Lagrange multiplier of the budget constraint in a related problem. In particular, suppose the planner seeks to maximize weighted consumer utility

$$\max_{(q,t)} \int_{\Theta} [\omega(\theta)v(q(\theta)) - t(\theta)] dF(\theta),$$

subject to (IC), (LB), (IR $_{\theta}$), (NLS), and the additional constraint

$$\int_{\Theta} [cq(\theta) - t(\theta)] dF(\theta) \leq B,$$

where B is the social planner's total budget for subsidy spending. In that case, Lagrangian duality implies there exists a Lagrange multiplier α^* such that the optimizer of that budget-constrained program is the same as the optimizer of (OPT) given weight α^* .

A microfoundation for the welfare weights $\omega(\theta)$ was offered in a redistributive mechanism design model by [Pai and Strack \(2024\)](#) as the expected marginal value for money of a consumer with concave preferences over money and can be adapted to this setting. In particular, suppose that consumers have two dimensions of private information I and θ , with a concave utility function

$$U(\theta) = \varphi(\theta v(q(\theta))) + I - t(\theta),$$

for a strictly concave and twice-differentiable $\varphi : \mathbb{R} \rightarrow \mathbb{R}$. Suppose that the social planner solves

$$\max_{(q,t)} \int_{\Theta} [\omega(\theta)U(\theta) - \alpha[cq(\theta) - t(\theta)]] dF(\theta),$$

subject to (IC), (LB), (IR $_{\theta}$), and (NLS). Then by the same argument as offered by [Pai and Strack \(2024\)](#), the optimal mechanism (q^*, t^*) in that problem is also the optimal mechanism of the program (OPT) with endogenously-determined welfare weights $\omega^*(\theta) = \mathbf{E}_I[\omega(\theta)\varphi'(\theta v(q(\theta))) + I - t(\theta) \mid \theta]$.

7 Discussion

In this section, we discuss the implications and limitations of our findings for subsidy programs observed in practice. Specifically, we consider applications of our results to three markets subsidized by various governments worldwide: food assistance programs, public transit, and pharmaceuticals. We also discuss the implications of our results for policy decisions not explicitly incorporated in the model, including the choice of product for subsidization, eligibility criteria and ordeals.

7.1 Implications of Our Results for Practical Subsidy Design

Food Assistance Programs. In the United States in 2021, the Supplemental Nutrition Assistance Program (SNAP, more commonly known as food stamps), provided \$111 billion in food assistance to over 41 million Americans ([Center on Budget and Policy Priorities, 2021](#)). The program primarily targets families with children, older adults, and people with disabilities, with eligibility generally requiring a household to have a gross income below 130% of the federal poverty line.¹⁴ SNAP benefits are capped monthly and delivered via Electronic Benefit Transfer (EBT) cards, which function like debit cards and can be used at authorized retailers to purchase food.

Even among consumers eligible for SNAP, the demand for food (and especially higher-quality food items) is likely to be positively correlated with socioeconomic status. Applying our model with decreasing welfare weights, we would expect the optimal subsidy mechanism for food consumption to be structured so that subsidies are available for consumption up to a certain level of need, beyond which additional consumption is unsubsidized. This is precisely how food stamps are designed: they cover a fixed amount of spending on food and are designed to allow—if not expect—recipients to top up their subsidized consumption in the private market.

Food assistance programs are also widely offered in developing economies, with an estimated 1.5 billion people worldwide covered by food subsidies, including large programs in India, Indonesia, Egypt, and Sri Lanka ([Alderman et al., 2017](#)). Often, these programs cover basic staples, like rice, coarse bread and cassava, consumed disproportionately by the poor ([Mackenzie, 1991](#)). In those cases, our model with positive correlation between demand and need may be most relevant, and our results predict that subsidies are sufficiently generous that consumers do not supplement their

¹⁴ At times, the eligibility rules and subsidy levels for food stamps have been a subject of debate. For example, in 2013, the program was criticized for changes that led to *more* qualifying families receiving *less* assistance, see [Severson and Hu \(2013\)](#).

consumption of the good in the private market.

Public Transit Fare Caps. Public transit subsidies are a common policy tool used by cities and countries around the world to promote affordable access to transportation. Demand for public transit often varies by demographic group. In the United States, for example, while about 5% of all workers reported commuting via public transportation in the 2019 American Community Survey, these commuters were more likely to be women, younger workers, and people of color (Burrows, Burd and McKenzie, 2021). Outside of seven “transit-heavy” metropolitan areas identified by Burrows et al. (2021) (including New York, San Francisco, and Washington D.C.), lower-income individuals were also much more likely to rely on public transportation, with 44.4% of transit commuters earning less than \$25,000 annually.¹⁵

Applying our results to this context,¹⁶ if the social planner places more weight on demographic groups relying heavily on public transportation, we would expect higher subsidies for high-demand users of public transportation, including more significant discounts for consumers who consume public transit more intensively. In practice, various cities offer transit subsidies toward targeted demographic groups, including New York City’s “Fair Fares” program, which offers half-price rides for lower-income New Yorkers (City of New York, 2024). Many public transportation authorities also offer “fare capping,” where public transit is free to consumers after a certain number of paid trips or level of spending within a day, week or month.¹⁷ These fare capping programs are consistent with the volume discounts in the optimal subsidy mechanism in the positive correlation case, identified in Section 5.¹⁸

Pharmaceutical Subsidies Some public health programs include pharmaceutical subsidies to ensure affordable access to medicines, particularly for vulnerable populations. To apply our model to the pharmaceutical market, we interpret the welfare weights as representing *social priorities*,

¹⁵ In larger cities, however, Burrows et al. (2021) observe that there is also a notable share of higher-income individuals using public transit, illustrating complexity in transit usage patterns across income groups.

¹⁶ Because public transport authorities can typically control the price of the public transit tickets used by consumers to top up subsidized consumption, our model including costly taxation in Section 6.2 may be most relevant here. The price of topping up tickets may be constrained in practice by the requirement that public transit tickets be available for tourists and occasional riders without enrolling in the subsidy program or by the social planner’s concern for inefficiencies created for consumers ineligible for the subsidy program.

¹⁷ See, e.g., the fare capping schemes in New York (Metropolitan Transportation Authority, 2024), London (Transport for London, 2024), Sydney (Transport for NSW, 2024), and Hong Kong (Hong Kong Special Administrative Region Government, 2024).

¹⁸ We note that zero marginal prices in the payment schedule (as in fare capping) arise in the limit as $q \rightarrow \infty$ of the optimal subsidy schedule when F has decreasing hazard rate and ω is increasing in θ .

which may place a greater weight on the sick, elderly, or disabled. Since the demand for healthcare and medicines is likely correlated with these social priorities, our results for positive correlation are most relevant and suggest that the social planner would offer greater subsidies to individuals with a high demand for pharmaceuticals.

These insights are reflected in the design of various pharmaceutical subsidy programs worldwide. For example, in the United States, both Medicaid and Medicare Part D provide pharmaceutical coverage for lower-income individuals and seniors, who tend to have greater demand for medicine ([Centers for Medicare & Medicaid Services, 2024a,b](#)). Similarly, Australia’s Pharmaceutical Benefits Scheme (PBS) requires all consumers to pay out-of-pocket fees for pharmaceuticals but offers reduced fees for vulnerable groups, such as pensioners and welfare recipients. The Australian PBS also includes a “Safety Net,” capping total out-of-pocket expenditure on pharmaceuticals per annum, after which additional pharmaceutical purchases are free ([Australian Government Department of Health and Aged Care, 2024](#)). Other countries, including New Zealand, Sweden, and Norway, have similar pharmaceutical spending caps.¹⁹ These caps are reminiscent of the optimal mechanism derived in Section 5 for positive correlation between demand and need.

7.2 Implications for Product Choice and Eligibility Restrictions

In our model, we treat the product subsidized and the set of eligible consumers as outside the social planner’s control. However, in some real-world settings, the subsidy designer can choose which products to subsidize and to whom subsidies are offered, based on observable consumer characteristics. Although these decisions are not explicitly integrated into our model, our results in this paper have several implications for such choices.

First, when welfare weights decrease in demand, which we expect to be a relevant assumption for many goods, the social planner only provides subsidies in the optimal mechanism if the average welfare weight exceeds the opportunity cost of funds. As a result, for such goods, restrictive eligibility criteria may be necessary to justify in-kind subsidy programs. This may explain why many subsidy programs observed in practice restrict eligibility for in-kind subsidies based on observable characteristics, including income, family status, citizenship, and other factors. For example, the U.S. government has tightened eligibility requirements for SNAP multiple times over

¹⁹ See, e.g., New Zealand’s Prescription Subsidy Scheme ([New Zealand Government, 2024](#)), Sweden’s “high-cost protection” (högkostnadsskydd) system ([Nordic Council of Ministers, 2024](#)), and Norway’s Frikort ([Helsenorge, 2024](#)).

the program’s history.²⁰ As another example, the Indonesian government has recently restricted fuel subsidies to ride-share operators and owners of vehicles with smaller engines ([The Jakarta Post, 2024](#)).

In contrast, when welfare weights increase with demand, the scope for redistribution is larger, requiring only that the *maximum* welfare weight exceeds the opportunity cost of funds. This reduces the need for explicit eligibility restrictions: instead, in-kind subsidy programs are *self-targeting*, with the greatest benefits of the subsidy program accruing to the consumers who purchase the largest quantities of the good. This may explain why public transit fare caps, discussed above, typically lack stringent eligibility criteria.

Our results also inform product choice for subsidy designers: given a selection of products with similar demand patterns, the social planner prefers to implement subsidies for goods with higher correlation between demand and income. Extending beyond our model, this might involve trading off a reduction in the desirability of the good against the ease of self-targeting the subsidy.

An example of this effect can be seen in the Egyptian Bread Subsidy Program, which subsidizes *baladi* bread, a staple food in Egypt ([Adams, 2000](#)). Over time, the Egyptian government reduced the quality of the subsidized bread by limiting the weight of the loaves and the type of wheat used. [Adams \(2000\)](#) found that these changes may have been self-targeting, as poorer consumers disproportionately consumed the subsidized coarse wheat bread, while wealthier consumers opted for higher-quality unsubsidized alternatives.

Another way to restrict eligibility is through the imposition of “ordeals,” which are non-monetary costs of accessing subsidy programs, including time-consuming application processes and waiting lists. Ordeals can be effective when high welfare weight consumers face lower ordeal costs (see, e.g., [Nichols and Zeckhauser, 1982](#), [Finkelstein and Notowidigdo \(2019\)](#)). Our results suggest that ordeals may also improve targeting by changing the *joint* distribution of welfare weights and demand. For example, if high-welfare weight consumers tend to have low ordeal costs independently of their demand for the good, this could improve the selection of eligible consumers, benefiting the social planner and the average eligible consumer.

8 Conclusion

Real-world subsidy programs often coexist with private markets for the same good. In this paper, we show that a consumer’s ability to access private markets for the good limits the redistributive

²⁰ For example, the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 and the Fiscal Responsibility Act of 2023 both imposed stricter work requirements for SNAP recipients without children.

power of the subsidy program compared to when the social planner can design the entire market for the good. When consumers have access to a private market, total subsidies paid must be increasing in demand for the good, making subsidies a better tool for redistribution when the social planner wants to redistribute toward higher-demand consumers than when they want to redistribute toward lower-demand consumers. This is likely why subsidy programs, in practice, often subsidize goods disproportionately consumed by less-advantaged consumers and why subsidy programs for other types of goods typically have narrow eligibility requirements.

While we have focused on subsidies designed by governments in the presence of private markets, there is, in principle, no reason that the principal in the subsidy design problem we have studied need be interpreted as a government agency or the outside option as a competitive market. For example, our methods could also be used to study subsidies offered by firms with market power in the presence of a competitive fringe aiming to win market share from rivals. Similar methods may also be relevant to analyze competing principals in a mechanism design setting. Expanding the methods in this paper along these lines is a promising direction for future research.

References

- ADAMS, R. H., JR (2000): “Self-targeted subsidies: The political and distributional impact of the Egyptian food subsidy system,” *Economic Development and Cultural Change*, 49, 115–136.
- AKBARPOUR, M. (R) E. BUDISH (R) P. DWORCZAK (R) S. D. KOMINERS (2024a): “An economic framework for vaccine prioritization,” *The Quarterly Journal of Economics*, 139, 359–417.
- AKBARPOUR, M. (R) P. DWORCZAK (R) S. D. KOMINERS (2024b): “Redistributive Allocation Mechanisms,” *Journal of Political Economy*, 132, 1831–1875.
- ALDERMAN, H., U. GENTILINI, AND R. YEMTSOV (2017): *The 1.5 Billion People Question: Food, Vouchers, or Cash Transfers?*: World Bank Publications.
- AMADOR, M., AND K. BAGWELL (2013): “The Theory of Optimal Delegation With an Application to Tariff Caps,” *Econometrica*, 81, 1541–1599.
- AMADOR, M., I. WERNING, AND G.-M. ANGELETOS (2006): “Commitment vs. Flexibility,” *Econometrica*, 74, 365–396.
- ATKINSON, A. B., AND J. E. STIGLITZ (1976): “The Design of Tax Structure: Direct Versus Indirect Taxation,” *Journal of Public Economics*, 6, 55–75.

- (2015): *Lectures on Public Economics*: Princeton University Press.
- AUSTRALIAN GOVERNMENT DEPARTMENT OF HEALTH AND AGED CARE (2024): “About the PBS,” <https://pbs.gov.au/info/about-the-pbs>, Accessed: 2024-10-04.
- BARLOW, R. E., D. J. BARTHOLOMEW, J. M. BREMNER, AND H. D. BRUNK (1972): *Statistical Inference Under Order Restrictions: The Theory and Application of Isotonic Regression*, Wiley Series in Probability and Mathematical Statistics, New York: Wiley.
- BARRON, E. N. (1983): “Remarks on a Simple Optimal Control Problem with Monotone Control Functions,” *Journal of Optimization Theory and Applications*, 41, 573–586.
- 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.
- BLOMQUIST, S., AND V. CHRISTIANSEN (1998): “Topping up or Opting Out? The Optimal Design of Public Provision Schemes,” *International Economic Review*, 39, 399–411.
- BURROWS, M., C. BURD, AND B. MCKENZIE (2021): “Commuting by Public Transportation in the United States: 2019,” April, <https://www.census.gov/content/dam/Census/library/publications/2021/acs/acs-48.pdf>.
- CALZOLARI, G., AND V. DENICOLÒ (2015): “Exclusive Contracts and Market Dominance,” *American Economic Review*, 105, 3321–51.
- CARROLL, G., AND I. SEGAL (2019): “Robustly Optimal Auctions with Unknown Resale Opportunities,” *Review of Economic Studies*, 86, 1527–1555.
- CENTER ON BUDGET AND POLICY PRIORITIES (2021): “The Housing Choice Voucher Program,” *Policy Basics*, available at <https://www.cbpp.org/sites/default/files/atoms/files/PolicyBasics-housing-1-25-13vouch.pdf>.
- CENTERS FOR MEDICARE & MEDICAID SERVICES (2024a): “Prescription Drugs | Medicaid,” <https://www.medicaid.gov/medicaid/prescription-drugs/index.html>, Accessed: 2024-10-04.
- (2024b): “What Medicare Part D Drug Plans Cover,” <https://www.medicare.gov/drug-coverage-part-d/what-medicare-part-d-drug-plans-cover>, Accessed: 2024-10-04.

- CHE, Y.-K., I. GALE, AND J. KIM (2013): “Assigning Resources to Budget-Constrained Agents,” *Review of Economic Studies*, 80, 73–107.
- CITY OF NEW YORK (2024): “Fair Fares NYC,” <https://www.nyc.gov/site/fairfares/index.page>, Accessed: 2024-10-04.
- CONDORELLI, D. (2013): “Market and Non-Market Mechanisms for the Optimal Allocation of Scarce Resources,” *Games and Economic Behavior*, 82, 582–591.
- CORRAO, R., J. P. FLYNN, AND K. A. SASTRY (2023): “Nonlinear Pricing with Underutilization: A Theory of Multi-part Tariffs,” *American Economic Review*, 113, 836–860.
- CURRIE, J. (1994): “Welfare and the Well-Being of Children: The Relative Effectiveness of Cash and In-Kind Transfers,” *Tax policy and the economy*, 8, 1–43.
- 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. (1975): “A Many-Person Ramsey Tax Rule,” *Journal of Public Economics*, 4, 335–342.
- DIAMOND, P. A., AND J. A. MIRRLEES (1972): “Optimal Taxation and Public Production: I,” *The American Economic Review*, 62, 238–238.
- DOLIGALSKI, P., P. DWORCZAK, J. KRZYSTA, AND F. TOKARSKI (2023): “Incentive Separability,” *Working paper*.
- DOWNES, A. (1957): *An Economic Theory of Democracy*: Harper & Row.
- DWORCZAK, P. (2020): “Mechanism Design With Aftermarkets: Cutoff Mechanisms,” *Econometrica*, 88, 2629–2661.
- DWORCZAK, P. (c) S. D. KOMINERS (c) M. AKBARPOUR (2021): “Redistribution Through Markets,” *Econometrica*, 89, 1665–1698.
- EDLIN, A. S., AND C. SHANNON (1998): “Strict monotonicity in comparative statics,” *Journal of Economic Theory*, 81, 201–219.
- FINKELSTEIN, A., AND M. J. NOTOWIDIGDO (2019): “Take-up and Targeting: Experimental Evidence from SNAP,” *The Quarterly Journal of Economics*, 134, 1505–1556.
- 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.
- GUESNERIE, R. (1981): *On taxation and incentives: Further reflections on the limits to redistribution*: Inst. für Ges.-und Wirtschaftswiss., Wirtschaftstheoretische Abt., Univ.
- HAMMOND, P. J. (1979): “Straightforward individual incentive compatibility in large economies,” *The Review of Economic Studies*, 46, 263–282.
- HARDY, G., J. LITTLEWOOD, AND G. PÓLYA (1934): *Inequalities*: Cambridge University Press.
- HEALTHCARE.GOV (2024): “Save on out-of-pocket costs with a Silver plan,” <https://www.healthcare.gov/lower-costs/save-on-out-of-pocket-costs/>, Accessed: 2024-10-20.
- HELSENGORGE (2024): “Exemption Card for Public Health Services,” <https://www.helsenorge.no/en/payment-for-health-services/exemption-card-for-public-health-services/>, Accessed: 2024-10-04.
- HONG KONG SPECIAL ADMINISTRATIVE REGION GOVERNMENT (2024): “Public Transport Fare Subsidy Scheme - Hong Kong,” <https://subsidy-enquiry.ptfss.hk/eng/index>, Accessed: 2024-10-04.
- JULLIEN, B. (2000): “Participation Constraints in Adverse Selection Models,” *Journal of Economic Theory*, 93, 1–47.
- KANG, M., AND C. Z. ZHENG (2020): “Pareto Optimality of Allocating the Bad,” *Working paper*.
- KANG, Z. Y. (2023): “The Public Option and Optimal Redistribution,” *Working paper*.
- (2024): “Optimal Indirect Regulation of Externalities,” *Working paper*.
- KANG, Z. Y., AND E. V. MUIR (2022): “Contracting and Vertical Control by a Dominant Platform,” *Working paper*.
- KANG, Z. Y., AND M. WATT (2024): “Optimal In-Kind Redistribution,” *Working paper*.
- KLEINER, A., B. MOLDOVANU, AND P. STRACK (2021): “Extreme Points and Majorization: Economic Applications,” *Econometrica*, 89, 1557–1593.
- LISCOW, Z., AND A. PERSHING (2022): “Why Is So Much Redistribution In-Kind and Not in Cash? Evidence From a Survey Experiment,” *National Tax Journal*, 75, 313–354.
- LOERTSCHER, S., AND E. V. MUIR (2022): “Monopoly Pricing, Optimal Rationing, and Resale,” *Journal of Political Economy*, 130, 566–635.

- MACKENZIE, G. A. (1991): “Price Subsidies,” in *Public Expenditure Handbook: A Guide to Public Policy Issues in Developing Countries* ed. by Chu, K.-Y., and Hemming, R. Washington, D.C.: International Monetary Fund, Chap. IX, 60–67.
- METROPOLITAN TRANSPORTATION AUTHORITY (2024): “OMNY Fare Capping,” <https://new.mta.info/fares/omny-fare-capping>, Accessed: 2024-10-04.
- MILGROM, P., AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70, 583–601.
- MILGROM, P., AND C. SHANNON (1994): “Monotone Comparative Statics,” *Econometrica*, 62, 157–180.
- MIRRELEES, J. A. (1976): “Optimal Tax Theory: A Synthesis,” *Journal of Public Economics*, 6, 327–358.
- (1986): “The Theory of Optimal Taxation,” *Handbook of Mathematical Economics*, 3, 1197–1249.
- 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.
- NEW ZEALAND GOVERNMENT (2024): “Prescription Subsidy Scheme,” <https://www.govt.nz/browse/health/gps-and-prescriptions/prescription-subsidy-scheme/>, Accessed: 2024-10-04.
- NICHOLS, A. L., AND R. J. ZECKHAUSER (1982): “Targeting Transfers through Restrictions on Recipients,” *American Economic Review*, 72, 372–377.
- NORDIC COUNCIL OF MINISTERS (2024): “Medicines and High-Cost Protection - Sweden,” <https://www.norden.org/en/info-norden/medicines-and-high-cost-protection-sweden>, Accessed: 2024-10-04.
- PAI, M., AND P. STRACK (2024): “Taxing Externalities Without Hurting the Poor,” *Working paper*.
- PALMER, J. (2021): “Why China Is Cracking Down on Private Tutoring,” *Foreign Policy*, available at <https://foreignpolicy.com/2021/07/28/china-private-tutoring-education-regulation-crackdown/>.

- PHILIPPON, T., AND V. SKRETA (2012): “Optimal Interventions in Markets with Adverse Selection,” *American Economic Review*, 102, 1–28.
- RAMSEY, F. P. (1927): “A Contribution to the Theory of Taxation,” *The Economic Journal*, 37, 47–61.
- REID, W. T. (1968): “A Simple Optimal Control Problem Involving Approximation by Monotone Functions,” *Journal of Optimization Theory and Applications*, 2, 365–377.
- REUTER, M., AND C.-C. GROH (2020): “Mechanism Design for Unequal Societies,” *Working paper*.
- ROBERTSON, T., F. WRIGHT, AND R. DYKSTRA (1988): *Order Restricted Statistical Inference*, Probability and Statistics Series: Wiley.
- SAEZ, E. (2002): “The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes,” *Journal of Public Economics*, 83, 217–230.
- SCHMIDT-KASPAREK, U. (2023): “PKV wirbt nur verhalten um freiwillig GKV-Versicherte,” <https://www.versicherungsmagazin.de/rubriken/branche/pkv-wirbt-nur-verhalten-um-freiwillig-gkv-versicherte-3428579.html>, Accessed: 2024-10-20.
- SEVERSON, K., AND W. HU (2013): “Cut in Food Stamps Forces Hard Choices on Poor,” *The New York Times*, <https://www.nytimes.com/2013/11/08/us/cut-in-food-stamps-forces-hard-choices-on-poor.html>, Accessed: 2024-10-04.
- STIGLITZ, J. E. (1987): “Pareto Efficient and Optimal Taxation and the New New Welfare Economics,” *Handbook of Public Economics*, 2, 991–1042.
- THE JAKARTA POST (2024): “Analysis: Govt plans to limit consumption of subsidized fuels again,” <https://www.thejakartapost.com/opinion/2024/09/12/analysis-govt-plans-to-limit-consumption-of-subsidized-fuels-again.html>, Accessed: 2024-10-12.
- 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.
- TRANSPORT FOR LONDON (2024): “Tube and Rail Fares - Pay-as-you-go Caps,” <https://tfl.gov.uk/fares/find-fares/tube-and-rail-fares/pay-as-you-go-caps>, Accessed: 2024-10-04.

- TRANSPORT FOR NSW (2024): “Opal Benefits,” <https://transportnsw.info/tickets-opal/opal/opal-benefits>, Accessed: 2024-10-04.
- VAN EEDEN, C. (1956): “Maximum likelihood estimation of ordered probabilities,” *Proceedings Koninklijke Nederlandse Akademie van Wetenschappen A*, 444–455.
- WATT, M. (2022): “Strong Monotonicity and Perturbation-Proofness of Exchange Economies,” *Working paper*.
- 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.
- YANG, F. (2021): “Costly Multidimensional Screening,” *Working paper*.
- YANG, F., P. DWORCZAK, AND M. AKBARPOUR (2024): “Comparison of Screening Devices,” *Working Paper*.
- YANG, K. H., AND A. K. ZENTEFIS (2024): “Monotone Function Intervals: Theory and Applications,” *Working paper*.

A Derivation of the Optimal Mechanism

In this section, we derive Theorem 2, and a generalization that permits more complex lower-bound constraints. We first derive Lemma 1 and 2, and then reformulate the resulting subsidy design problem (OPT) as a convex program in q . We then discuss a generalization of that program that we study in this section. We then derive the optimal allocation rule for that generalized convex program and discuss our proof approach in the context of the mechanism design literature.

A.1 Proofs of Lemma 1 and Lemma 2

We first present a proof of Lemma 1, establishing the lower-bound constraint on the allocation rule.

Proof. For the “only if” direction, let q be an allocation rule, and suppose that (q, t) is implemented by a marginal subsidy function $\sigma : [0, A] \rightarrow \mathbb{R}_+$. Then, given payment schedule P^σ , a consumer of type θ solves the maximization problem

$$\max_{z \in [0, A]} \theta v(z) - P^\sigma(z),$$

which has the first-order condition

$$\theta v'(q(\theta)) - (c - \sigma(q(\theta))) = 0, \text{ or equivalently, } v'(q(\theta)) = \frac{c - \sigma(q(\theta))}{\theta}.$$

Since $\sigma(z) \geq 0$, we have $v'(q(\theta)) \leq v'(q^{\text{LF}}(\theta)) = c/\theta$, which implies $q(\theta) \geq q^{\text{LF}}(\theta)$ by the strict concavity of v . The monotonicity of q in θ follows by the supermodularity of the consumer’s optimization problem and the Topkis theorem (cf. Milgrom and Shannon, 1994).

For the “if” direction, suppose that q is nondecreasing in θ and satisfies $q(\theta) \geq q^{\text{LF}}(\theta)$. We construct a marginal subsidy rule $\sigma : [0, A] \rightarrow \mathbb{R}_+$ satisfying (SUB) such that $(q, P^\sigma \circ q)$ satisfies (IC). For that purpose, for any $z \in [0, A]$, define $q^{-1}(z)$ as the generalized inverse of q , so that $q^{-1}(z) = \inf\{\theta \in \Theta : q(\theta) \geq z\}$ or $\bar{\theta}$ if that set is empty. Note that q^{-1} is nondecreasing because q is nondecreasing. Let $\sigma(z) := c - q^{-1}(z)v'(z)$, so

$$P^\sigma(q) = \int_0^q q^{-1}(z)v'(z) dz.$$

Because q^{-1} is nondecreasing, the consumer’s objective $\theta v(z) - P^\sigma(z)$ is concave in z , and, by

construction, $q(\theta) \in \arg \max_z \theta v(z) - P^\sigma(z)$ for each $\theta \in \Theta$, so $(q, P^\sigma \circ q)$ satisfies (IC). On the other hand, (SUB) is satisfied because $\sigma(z) \leq c$ for all $z \in [0, A]$. \square

We now establish Lemma 2, which establishes the constraints on payments within the mechanism and on the feasible payoffs for consumers.

Proof. Suppose that (q, t) satisfies (IC) and (LB).

For the “only if” direction, suppose that (q, t) is implemented by a marginal subsidy function $\sigma : [0, A] \rightarrow \mathbb{R}_+$. Then because $P^\sigma(q) = cq - \Sigma(q)$, (SUB) implies $t(\theta) = P^\sigma(q(\theta)) \geq cq(\theta) - cq(\theta) = 0$, which is (NLS). On the other hand, because $P^\sigma(z) \leq cz$ for all $z \in [0, A]$, $U^{\text{LF}}(\underline{\theta}) \leq \underline{\theta}v(q^{\text{LF}}(\underline{\theta})) - P^\sigma(q^{\text{LF}}(\underline{\theta}))$. But because $q(\underline{\theta}) \in \arg \max_z \theta v(z) - P^\sigma(z)$, we have $\underline{U} = \underline{\theta}v(q(\underline{\theta})) - P^\sigma(q(\underline{\theta})) \geq \underline{\theta}v(q^{\text{LF}}(\underline{\theta})) - P^\sigma(q^{\text{LF}}(\underline{\theta})) \geq U^{\text{LF}}(\underline{\theta})$, implying (IR $_{\underline{\theta}}$).

The converse is obtained using the same construction of σ presented in the proof of Lemma 1. \square

A.2 Reformulating the Subsidy Design Problem as a Convex Program

Proposition 14 (reformulating the subsidy design problem). *The subsidy design problem (OPT) is equivalent to the following convex program:*

$$\max_{q: \Theta \rightarrow [0, A]} \alpha \int_{\Theta} [J(\theta)v(q(\theta)) - cq(\theta)] \, dF(\theta) + (\text{terms independent of } q), \quad (\text{OPT-}q)$$

such that q is nondecreasing and satisfies for all $\theta \in \Theta$, $q(\theta) \geq q^{\text{LF}}(\theta)$,

where $J(\theta)$ is the virtual welfare, defined by

$$J(\theta) = \theta + \frac{\int_{\theta}^{\bar{\theta}} [\omega(s) - \alpha] \, dF(s) + \max\{\mathbf{E}[\omega] - \alpha, 0\} \cdot \theta \delta_{\underline{\theta}}(\theta)}{\alpha f(\theta)},$$

and $\delta_{\underline{\theta}}(\theta)$ is the Dirac delta (point mass) centered at $\underline{\theta}$.

Proof. We apply standard mechanism design techniques to eliminate various constraints and express (OPT) in terms of the allocation rule q .

First, we use Myerson’s (1981) Lemma to replace the (IC) constraint with the requirement that q be a nondecreasing function of θ , writing \mathcal{Q} for the set of nondecreasing functions on Θ .

Second, given the (IC) and (LB) constraints, it suffices to enforce the individual rationality constraint only for the lowest type, that is, to ensure that $\underline{U} := \underline{\theta}v(q(\underline{\theta})) - t(\underline{\theta}) \geq U^{\text{LF}}(\underline{\theta})$. To see

that, note that (LB) implies $v(q(\theta)) \geq v(q^{\text{LF}}(\theta))$ because v is increasing, which when combined with $\underline{U} \geq U^{\text{LF}}(\underline{\theta})$ implies that

$$\underline{U} + \int_{\underline{\theta}}^{\theta} v(q(\theta)) \, d\theta \geq U^{\text{LF}}(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} v(q^{\text{LF}}(\theta)) \, d\theta,$$

so $U(\theta) \geq U^{\text{LF}}(\theta)$ by the Milgrom and Segal (2002) envelope theorem.

Third, we again apply the envelope theorem to express payments in terms of the allocation rule, allowing us to rewrite the social planner's optimization program as

$$\max_{(q,t)} \left\{ (\mathbf{E}[\omega] - \alpha) \underline{U} + \alpha \int_{\Theta} \left[\left(\theta + \frac{\int_{\underline{\theta}}^{\bar{\theta}} [\omega(s) - \alpha] \, dF(s)}{\alpha f(\theta)} \right) v(q(\theta)) - cq(\theta) \right] dF(\theta) \right\},$$

such that $q \in \mathcal{Q}$, $\underline{U} \geq U^{\text{LF}}(\underline{\theta})$, and (q, t) satisfies (NLS) and (LB).

Fourth, we argue that—given the (IC) constraint—it suffices to enforce (NLS) only for the lowest type $\underline{\theta}$. To see this, note that for any two types $\theta, \theta' \in \Theta$,

$$\begin{aligned} t(\theta') - t(\theta) &= \left[\theta' v(q(\theta')) - U(\underline{\theta}) - \int_{\underline{\theta}}^{\theta'} v(q(s)) \, ds \right] - \left[\theta v(q(\theta)) - U(\underline{\theta}) - \int_{\underline{\theta}}^{\theta} v(q(s)) \, ds \right] \\ &= \theta' v(q(\theta')) - \theta v(q(\theta)) - \int_{\theta}^{\theta'} v(q(s)) \, ds \\ &= \int_{\theta}^{\theta'} s v'(q(s)) \, dq(s), \end{aligned}$$

where the last equality employs integration-by-parts for the Lebesgue-Stieltjes integral. This expression is nonnegative for $\theta \geq \theta'$, implying that the total payment is increasing in θ . As a result, if $t(\underline{\theta})$ is nonnegative, so is the payment made by all higher types.

Finally, fixing q , we note that when the opportunity cost of funds is low (so $\mathbf{E}[\omega] > \alpha$), the social planner prefers to implement q using a payment rule t that makes the lowest type's utility \underline{U} as large as possible (making the additive term outside the integrand in the objective above as large as possible), which entails choosing $t(\underline{\theta}) = 0$, so the (NLS) constraint binds. Because $\underline{U} = \underline{\theta} v(q(\underline{\theta}))$ depends on the choice of q , we incorporate it into the integrand of the objective as a point mass at $\underline{\theta}$. Conversely, when the cost of funds is high (so $\mathbf{E}[\omega] < \alpha$), the social planner prefers a payment rule t making \underline{U} as small as possible, which entails choosing $t(\underline{\theta})$ such that the (IR $_{\underline{\theta}}$) constraint binds and $\underline{U} = U^{\text{LF}}(\underline{\theta})$. Finally, if $\mathbf{E}[\omega] = \alpha$, the social planner is indifferent over

transfer rules implementing q .

These findings allow us to write the mechanism design problem as follows:

- (a) If the opportunity cost of funds is low, so $\mathbf{E}[\omega] > \alpha$:

$$\max_{q \in \mathcal{Q}: q \geq q^{\text{LF}}} \alpha \int_{\Theta} [J(\theta)v(q(\theta)) - cq(\theta)] dF(\theta), \quad (\text{OPT-}\ell)$$

where

$$J(\theta) = \theta + \frac{\int_{\theta}^{\bar{\theta}} [\omega(s) - \alpha] dF(s) + (\mathbf{E}[\omega(\theta)] - \alpha)\underline{\theta}\delta_{\underline{\theta}}}{\alpha f(\theta)}.$$

- (b) If the opportunity cost of funds is high, so $\mathbf{E}[\omega] \leq \alpha$:

$$\max_{q \in \mathcal{Q}: q \geq q^{\text{LF}}} \alpha \int_{\Theta} [J(\theta)v(q(\theta)) - cq(\theta)] dF(\theta) + (\mathbf{E}[\omega] - \alpha)\underline{\theta}v(q^{\text{LF}}(\underline{\theta})), \quad (\text{OPT-}h)$$

where

$$J(\theta) = \theta + \frac{\int_{\theta}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)}{\alpha f(\theta)}.$$

These two expressions for virtual welfare differ by a point mass at $\underline{\theta}$ included in the expression for virtual welfare when the opportunity cost of funds is low (corresponding to the (NLS) constraint binding for $\underline{\theta}$, so the designer wishes to allocate some quantity of the good to all types for free). The objectives otherwise differ only in an additive term independent of q that appears in (OPT- h).

Putting these expressions together leads to the convex program in the statement of Proposition 14. \square

A.3 Generalizing (OPT- q) and Theorem 2

In this subsection, we introduce a generalization of the (LB) constraint that requires q to exceed pointwise an arbitrary continuous $q^{\text{L}} \in \mathcal{Q}$.²¹ For this purpose, in order to ensure that it is possible to induce a consumer to consume at least q^{L} units of the good, we assume in this section an Inada condition on v : that $\lim_{z \rightarrow \infty} v'(z) = 0$ and $\lim_{z \rightarrow 0} v'(z) = +\infty$.

Putting these assumptions together, we study the more general problem than (OPT- q), given

²¹ In fact, the restriction to \mathcal{Q} is without loss of generality, because even if $q^{\text{L}} \notin \mathcal{Q}$, any feasible q must be pointwise greater than the monotone envelope of q^{L} (the least monotone function pointwise greater than q^{L}).

by

$$\max_{\substack{q: \Theta \rightarrow [0, A] \text{ s.t.} \\ q \text{ is non-decreasing}}} \int_{\Theta} [J(\theta)v(q(\theta)) - cq(\theta)] dF(\theta) \text{ subject to } \forall \theta \in \Theta, \text{ and } q(\theta) \geq q^L(\theta). \quad (\mathbf{P})$$

In this section, we write $D(\theta)$ for the demand of type θ (suppressing the price c , which is held fixed in this section), and let $D^{-1}(q)$ be its inverse, i.e., $D^{-1}(q) = \inf\{\theta \in \Theta | d \geq q \text{ for } d \in D(\theta)\}$, which is well-defined by the strict monotonicity of v' and the Inada condition. We write $L(\theta) = D^{-1}(q^L(\theta))$, which is the type demanding $q^L(\theta)$ units of the good in the laissez-faire economy at price c . We note that L is monotone and continuous by the monotonicity and continuity of q^L and D (and therefore D^{-1}).

We now state the generalization of Theorem 2 characterizing the solution to (P).

Theorem 2'. *The solution to (P), $q^* : \Theta \rightarrow [0, A]$, satisfies*

$$q^*(\theta) = D(H(\theta))$$

where

$$H(\theta) = \begin{cases} L(\theta) & \text{if } L(\theta) \geq \overline{J}_{[\underline{\theta}, \theta]}(\theta), \\ \overline{J}_{[\underline{\theta}, \kappa_+(\theta)]}(\theta) & \text{otherwise,} \end{cases}$$

and

$$\kappa_+(\theta) = \inf \left\{ \hat{\theta} \in \Theta : \hat{\theta} \geq \theta, \text{ and } L(\hat{\theta}) \geq \overline{J}_{[\underline{\theta}, \hat{\theta}]}(\hat{\theta}) \right\} \text{ or } \bar{\theta} \text{ if that set is empty.}$$

Finally, if L is continuous on $\text{int } \Theta$, then q^* is continuous as well.

Note that Theorem 2 follows as its consequence because if $q^L(\theta) = q^{\text{LF}}(\theta)$, the function $L(\theta)$ defined above is just θ .

Proof. We break the proof of Theorem 2' into a number of steps. We first present some preliminary lemmata establishing properties of the ironing operator. We then use those lemmata to establish the equivalence of the various expressions for H in Theorem 2', and the feasibility of the expression for q^* in (P). Finally, we establish the optimality of q^* in (P) using variational inequality arguments.

A.4 Preliminary Ironing Lemmata.

To state our preliminary lemmata on ironing, we require additional notation. Define the F -weighted average of ϕ on $[\theta_1, \theta_2]$ with $\theta_1 \leq \theta_2$ as follows

$$\mu(\theta_1, \theta_2) = \begin{cases} \frac{\int_{\theta_1}^{\theta_2} \phi(s) dF(s)}{F(\theta_2) - F(\theta_1)} & \text{if } \theta_2 > \theta_1 \\ \phi(\theta_1) & \text{if } \theta_2 = \theta_1. \end{cases}$$

Recall that on any ironing interval $[\theta_1, \theta_2]$ of ϕ , for all $\theta \in [\theta_1, \theta_2]$, we have $\bar{\phi}(\theta) = \mu(\theta_1, \theta_2)$. In this section, we will include the singleton $[\theta_1, \theta_1]$ in our definition of an ironing interval, for which this expression is also valid.

The first lemma describes the effect of “interrupting” an ironing interval—by which we mean restricting attention to a subinterval—on the F -weighted average.

Lemma 4 (Interrupted Ironing). *Let $\bar{\phi}(\theta) = c$ and let $[\theta_1^*, \theta_2^*]$ be an ironing interval containing θ , so $\mu(\theta_1^*, \theta_2^*) = c$. Then for any $\theta_1 \in [\theta_1^*, \theta_2^*]$, we have $\mu(\theta_1, \theta_2^*) \leq c$ and for any $\theta_2 \in [\theta_1^*, \theta_2^*]$, we have $\mu(\theta_1^*, \theta_2) \geq c$.*

Proof. This follows by definition of ironing as the convex envelope because $\int_{\theta_1^*}^{\theta_2} \phi(s) dF(s) \geq \int_{\theta_1^*}^{\theta_2} c dF(s)$, so $\mu(\theta_1^*, \theta_2) \geq c$. Similarly, since $\int_{\theta_1}^{\theta_1^*} \phi(s) dF(s) \geq \int_{\theta_1}^{\theta_1^*} c dF(s)$ and $\int_{\theta_1}^{\theta_2^*} \phi(s) dF(s) = \int_{\theta_1}^{\theta_1^*} \phi(s) dF(s) + \int_{\theta_1^*}^{\theta_2^*} \phi(s) dF(s) \geq \int_{\theta_1}^{\theta_1^*} c dF(s) + \int_{\theta_1^*}^{\theta_2^*} c dF(s) = \int_{\theta_1}^{\theta_2^*} c dF(s)$, we have $\int_{\theta_1}^{\theta_2^*} \phi(s) dF(s) \leq \int_{\theta_1}^{\theta_2^*} c dF(s)$, which implies $\mu(\theta_1, \theta_2^*) \leq c$. \square

We use this result to establish the following useful characterization of the ironing operator (cf. Barlow et al., 1972; Robertson et al., 1988, who obtain a similar formula for a discrete ironing operator arising in the study of isotonic regression).

Lemma 5 (max-min formula for ironing). *For any $\phi : \Theta \rightarrow \mathbb{R}$ and $\theta \in [\theta_\ell, \theta_h] \subseteq \Theta$, we have*

$$\overline{\phi|_{[\theta_\ell, \theta_h]}}(\theta) = \max_{\theta_1: \theta_\ell \leq \theta_1 \leq \theta} \min_{\theta_2: \theta \leq \theta_2 \leq \theta_h} \mu(\theta_1, \theta_2).$$

Moreover, writing θ_1^* and θ_2^* for the solutions to that minimax problem, then $[\theta_1^*, \theta_2^*]$ is a (possibly singleton) ironing interval for $\phi|_{[\theta_\ell, \theta_h]}$ containing θ .

Proof. Let $\bar{\phi}(\theta) = c$ and let $\hat{\theta}_1$ and $\hat{\theta}_2$ be the endpoints of the maximal ironing interval containing θ .

We first show that $\theta_1^* \geq \hat{\theta}_1$ and $\theta_2^* \leq \hat{\theta}_2$. To see this, fix any $\theta_1 \leq \theta$, and suppose (for a contradiction) that $\theta_2^* > \hat{\theta}_2$. We show that $\mu(\theta_1, \hat{\theta}_2) < \mu(\theta_1, \theta_2^*)$, contradicting the optimality of θ_2^* .

By the monotonicity of $\bar{\phi}$, we have $\bar{\phi}(\theta_2^*) = c' > c$. If θ_2^* is inside an ironing interval of ϕ , Lemma 4 implies that we may as well assume that θ_2^* is the right endpoint of that ironing interval, because that reduces μ on the ironing interval and thus also on (θ_1, θ_2^*) . But then by the monotonicity of $\bar{\phi}$, we can write $(\hat{\theta}_2, \theta_2^*]$ as the union of intervals on which $\mu > c$. But then removing the contribution of each such interval would lower the value of μ , so $\mu(\theta_1, \hat{\theta}_2) < \mu(\theta_1, \theta_2^*)$. This implies that $\theta_2^* \leq \hat{\theta}_2$. The argument for $\theta_1^* \geq \hat{\theta}_1$ is analogous.

We can thus restrict attention to decision variables θ_1 and θ_2 within $[\hat{\theta}_1, \hat{\theta}_2]$. Fix any θ_1 , then by Lemma 4, we must have that

$$\min_{\theta \leq \theta_2 \leq \hat{\theta}_2} \mu(\theta_1, \theta_2) \leq \mu(\theta_1, \hat{\theta}_2) \leq \mu(\hat{\theta}_1, \hat{\theta}_2) = c.$$

But then

$$\max_{\hat{\theta}_1 \leq \theta_1 \leq \theta} \min_{\theta \leq \theta_2 \leq \hat{\theta}_2} \mu(\theta_1, \theta_2) \leq c.$$

But this bound is obtained for *any* θ_1^*, θ_2^* that are the endpoints of an ironing interval containing θ , so that

$$\max_{\theta_1: \theta_\ell \leq \theta_1 \leq \theta} \min_{\theta_2: \theta \leq \theta_2 \leq \theta_h} \mu(\theta_1, \theta_2) = \overline{\phi|_{[\theta_\ell, \theta_h]}}(\theta),$$

as required.

On the other hand, if θ_1^* and θ_2^* are solutions to the minimax problem, we have $\mu(\theta_1^*, \theta_2^*) = \mu(\hat{\theta}_1, \hat{\theta}_2)$. But then by the construction of $\bar{\phi}$ (as the slope of the convex envelope of $\text{co}\Phi$), the interval $[\theta_1^*, \theta_2^*]$ is an ironing interval of $\bar{\phi}$ containing θ as well. \square

The expression for the ironing operator as the solution of a minimax problem obtained in Lemma 5 allows us to establish the monotonicity of the ironing operator in the domain of its argument:

Lemma 6 (domain expansions for ironing). *Let $\theta'_\ell < \theta_\ell < \theta_h < \theta'_h$, then for all $\theta \in [\theta_\ell, \theta_h]$, we have:*

- (a) $\overline{\phi|_{[\theta'_\ell, \theta_h]}}(\theta) \geq \overline{\phi|_{[\theta_\ell, \theta_h]}}(\theta)$, with strict inequality if and only if θ_ℓ is in the interior of an ironing interval of ϕ in $[\theta'_\ell, \theta_h]$.
- (b) $\overline{\phi|_{[\theta_\ell, \theta'_h]}}(\theta) \leq \overline{\phi|_{[\theta_\ell, \theta_h]}}(\theta)$, with strict inequality if and only if θ_h is in the interior of an ironing interval of ϕ in $[\theta_\ell, \theta'_h]$.

Lemma 6 states that expanding the ironing domain to the left can increase (but not decrease)

the value of the ironed function on the original domain, while expanding the ironing domain to the right can only decrease its value.

Proof. By the max-min expression for $\overline{\phi}_{|\theta_\ell, \theta_h]}$ in Lemma 5, it is clear that reducing θ_ℓ to $\theta'_\ell < \theta_\ell$ can only *increase* the value of the ironed function at any θ (by expanding the set being maximized over) and that increasing θ_h to $\theta'_h > \theta_h$ can only *decrease* the value (by expanding the set being minimized over). On the other hand, if the objective value decreases strictly, the new optimizers must be in the expanded choice set (e.g., if θ_ℓ decreases to θ'_ℓ , the optimizer for θ_1 must be in $(\theta'_\ell, \theta_\ell)$). But then, (θ_1^*, θ_2^*) must be a (non-singleton) ironing interval containing θ in its interior. \square

A.5 Equivalent Expressions for H and Feasibility of q^* .

In this section, we establish the equivalence of the various expressions for H given in Theorem 2 and the feasibility of the resulting q^* in (OPT- q).

Lemma 7 (equivalent representations of H and feasibility). *The subsidy type H is nondecreasing, satisfies for all $\theta \in \Theta$, $H(\theta) \geq L(\theta)$, and can equivalently be calculated as follows:*

(a)

$$H(\theta) = \begin{cases} L(\theta) & \text{if } L(\theta) \geq \overline{J}_{|\underline{\theta}, \theta]}(\theta), \\ \overline{J}_{|\kappa_-(\theta), \kappa_+(\theta)]}(\theta) & \text{otherwise,} \end{cases},$$

where $\kappa_-(\theta) = \sup \left\{ \hat{\theta} \in \Theta : \hat{\theta} \leq \theta, \text{ and } L(\hat{\theta}) \geq \overline{J}_{|\underline{\theta}, \hat{\theta}] }(\hat{\theta}) \right\}$ or $\underline{\theta}$ if that set is empty.

(b) $H = \overline{H^T}$, where

$$H^T(\theta) = \begin{cases} L(\theta) & \text{if } L(\theta) \geq \overline{J}_{|\underline{\theta}, \theta]}(\theta), \\ J(\theta) & \text{otherwise.} \end{cases}$$

We introduce additional notation for the proof of Lemma 7. By the monotonicity of $L(\theta)$ and the continuity of $\overline{J}_{|\underline{\theta}, \theta]}(\theta)$, the type space Θ can be partitioned into alternating intervals within which either (a) $L(\theta) \geq \overline{J}_{|\underline{\theta}, \theta]}(\theta)$, or (b) $L(\theta) < \overline{J}_{|\underline{\theta}, \theta]}(\theta)$ in the interior of the interval. Call these *type (a) intervals* and *type (b) intervals*, respectively. Write these intervals (intersecting only at their endpoints) as $\Theta^1, \Theta^2, \dots$ so $\Theta = \cup_i \Theta^i$ and let $\text{int } \Theta^i = (\theta_-^i, \theta_+^i)$.

Let

$$\tilde{H}(\theta) = \begin{cases} L(\theta) & \text{for } \theta \text{ in any type (a) interval } \Theta^i \\ \overline{J}_{|\theta_-^i, \theta_+^i]}(\theta) & \text{for } \theta \text{ in any type (b) interval } \Theta^i, \end{cases}$$

matching the expression in part (a) of Lemma 7. To establish the *equivalence* of our expressions for H , we need to show that $\tilde{H} = H = \overline{H^T}$. Once we have proven those equivalent representations, to establish *feasibility* of q^* , we need only establish that at least one of \tilde{H} , H , or $\overline{H^T}$ are nondecreasing and bounded below by L : this will imply that $q^* = D(H(\cdot))$ is nondecreasing and bounded below by q^L .

Proof. We first show that $\tilde{H} = H$. This is clear for any θ in a type (a) interval, so we need only show for θ in type (b) intervals that $\tilde{H}(\theta) = \overline{J_{[\theta_-, \theta_+]}(\theta)} = \overline{J_{[\theta, \theta_+]}(\theta)} = H(\theta)$. By Lemma 6, it suffices to show that θ_-^i is not contained in an ironing interval in $[\theta, \theta_+^i]$ as this will imply $\overline{J_{[\theta_-, \theta_+]}(\theta)} = \overline{J_{[\theta, \theta_+]}(\theta)}$, as required. Suppose otherwise, then there must exist an ironing interval of $J_{[\theta, \theta_+]}(\theta)$ containing θ_-^i ending at some $\hat{\theta} \in \Theta^i$. Because $\hat{\theta} \in \Theta^i$, we must have that $\overline{J_{[\theta, \theta_+]}(\hat{\theta})} \geq L(\hat{\theta})$ which implies $\overline{J_{[\theta, \theta_+]}(\theta_-^i)} \geq L(\hat{\theta}) \geq L(\theta_-^i)$. But then for $\theta \leq \theta_-^i$ contained in the same ironing interval, we must have (by Lemma 6) that $\overline{J_{[\theta, \theta]}(\theta)} \geq \overline{J_{[\theta, \theta_+]}(\theta)} \geq L(\hat{\theta}) \geq L(\theta)$, which contradicts the fact that θ is in a type (a) interval.

We now show that \tilde{H} is continuous and nondecreasing in θ . These properties are immediate *within* any interval Θ^i , but to show them *across* adjacent intervals, it suffices to show that the two expressions for \tilde{H} match at the meeting points of adjacent intervals. But by construction, for any type (b) interval $[\theta_-^i, \theta_+^i]$, we have $\overline{J_{[\theta, \theta_-^i]}(\theta_-^i)} = L(\theta_-^i)$ and $\overline{J_{[\theta, \theta_+^i]}(\theta_+^i)} = L(\theta_+^i)$, and then, by our previous finding, we have that $\tilde{H}(\theta_-^i) = \overline{J_{[\theta_-, \theta_+]}(\theta_-^i)} = L(\theta_-^i)$, and $\tilde{H}(\theta_+^i) = \overline{J_{[\theta_-, \theta_+]}(\theta_+^i)} = L(\theta_+^i)$, as required.

Thus, \tilde{H} is nondecreasing and continuous, which implies that $\tilde{\mathcal{H}} := \int_{\underline{\theta}}^{\theta} \tilde{H}(s) dF(s)$ is a convex function. Our next goal is to show that \tilde{H} is the convex envelope of $\mathcal{H}^T := \int_{\underline{\theta}}^{\theta} H^T(s) dF(s)$, which by the uniqueness of the ironing $\overline{H^T}$ will imply that $\tilde{H} = H$.

To obtain that result, we first show that $\tilde{\mathcal{H}}$ is a minorant of \mathcal{H}^T . To see this, note that within any interval Θ^i , the function $\theta \mapsto \int_{\theta_-^i}^{\theta} \tilde{H}(s) dF(s)$ is the convex envelope of $\theta \mapsto \int_{\theta_-^i}^{\theta} H^T(s) dF(s)$. This means that for all $\theta \in \Theta^i$, $\int_{\theta_-^i}^{\theta} \tilde{H}(s) dF(s) \leq \int_{\theta_-^i}^{\theta} H^T(s) dF(s)$ with equality at $\theta = \theta_+^i$. But then, applying those facts from left to right, we obtain that $\tilde{\mathcal{H}}(\theta) \leq \mathcal{H}^T(\theta)$ for all $\theta \in \Theta$, and that $\tilde{\mathcal{H}}(\theta) = \mathcal{H}^T(\theta)$ for any θ at the endpoint of an interval. That implies that $\tilde{\mathcal{H}}$ is a minorant of \mathcal{H}^T on Θ .

Now suppose that there were a pointwise larger convex minorant of \mathcal{H}^T , say \mathcal{H}' (i.e., a convex function satisfying $\mathcal{H}^T \geq \mathcal{H}' \geq \mathcal{H}$ with $\mathcal{H}' \neq \mathcal{H}$). Since $\tilde{\mathcal{H}} = \mathcal{H}^T$ at the endpoint of each interval Θ^i , $\mathcal{H}' = \mathcal{H}^T$ there as well. So if $\mathcal{H}'(\theta) > \tilde{\mathcal{H}}(\theta)$ for some $\theta \in \Theta$, it must be for some $\theta \in \text{int}(\Theta^i)$. But then $\mathcal{H}'|_{\Theta^i}$ would be a convex minorant of $\mathcal{H}^T|_{\Theta^i}$ with $\mathcal{H}'|_{\Theta^i} > \tilde{\mathcal{H}}|_{\Theta^i}$, which would contradict the fact that $\tilde{\mathcal{H}}|_{\Theta^i}$ is the convex envelope of $\mathcal{H}^T|_{\Theta^i}$. This implies that $\tilde{\mathcal{H}}$ is the convex envelope of

\mathcal{H}^T , and so $\tilde{H} = H$.

Finally, having shown that $\tilde{H}(= H)$ is nondecreasing, for the feasibility of $q^* = D(c, H(\theta))$, we need only show that for all $\theta \in \Theta^i$, $\tilde{H}(\theta) \geq L(\theta)$. This is clearly satisfied on each type (a) interval, but within any type (b) interval, note by construction that $\overline{J|_{[\underline{\theta}, \theta]}}(\theta) \geq L(\theta)$. Since, by the above, we have $\overline{J|_{[\underline{\theta}, \theta]}}(\theta) = \overline{J|_{[\theta^i, \theta]}}(\theta)$, we have $\tilde{H}(\theta) \geq L(\theta)$ as well. \square

A.6 Optimality of q^* .

Having established that $q^*(\theta) = D(H(\theta), c)$ is feasible in (P), we now show that it is optimal.

Solving A Simpler Convex Program. To establish the optimality of q^* , we relate the solution of (OPT- q) to the solution of a simpler mechanism design problem, for which the optimum is the pointwise maximizer of the objective. In particular, we study the problem

$$\max_{q \in \mathcal{Q}} \int_{\Theta} [H(\theta)v(q(\theta)) - cq(\theta)] dF(\theta) \text{ subject to } \forall \theta \in \Theta, q(\theta) \geq q^L(\theta). \quad (\text{OPT-H})$$

By the above, we have that $q^*(\theta) = D(H(\theta))$ is feasible in (OPT-H), and, by the definition of D , it is also the pointwise maximizer of the objective in (OPT-H). Consequently, q^* solves (OPT-H).

Relating the Optimality Conditions. We use the optimality of q^* in (OPT-H) to demonstrate its optimality in (OPT- q). The optimality condition of (OPT- q) may be written as a variational inequality: for any other feasible q , we have that

$$\int_{\Theta} [J(\theta)v(q^*(\theta)) - cq^*(\theta)] dF(\theta) \geq \int_{\Theta} [J(\theta)v(q(\theta)) - cq(\theta)] dF(\theta). \quad (\text{VI})$$

Having already shown—via the optimality of q^* in (OPT-H)—that for any feasible q

$$\int_{\Theta} [H(\theta)v(q^*(\theta)) - cq^*(\theta)] dF(\theta) \geq \int_{\Theta} [H(\theta)v(q(\theta)) - cq(\theta)] dF(\theta),$$

it suffices to sign the difference in these inequalities:

$$\int_{\Theta} [J(\theta) - H(\theta)] [v(q^*(\theta)) - v(q(\theta))] dF(\theta) \geq 0.$$

Signing the Difference in the Variational Inequalities. To that end, we rewrite the previous inequality as

$$\int_{\Theta} [J(\theta) - H^{\text{T}}(\theta)] [v(q^*(\theta)) - v(q(\theta))] \, dF(\theta) + \int_{\Theta} [H^{\text{T}}(\theta) - H(\theta)] [v(q^*(\theta)) - v(q(\theta))] \, dF(\theta) \geq 0,$$

and sign the terms in that sum individually.

First, we establish that

$$\int_{\Theta} [H^{\text{T}}(\theta) - H(\theta)] v(q^*(\theta)) \, dF(\theta) = 0.$$

Intuitively, that follows because $H = \overline{H^{\text{T}}}$, and on any interval where $H \neq H^{\text{T}}$, H is the F -average of H^{T} , and q^* is constant, so $v(q^*(\theta))$ may be moved outside the integrand. Formally, consider any $c \in \text{im } H$, and let $\Theta_c := \{\theta \in \Theta : H(\theta) = c\}$. Because $H = \overline{H^{\text{T}}}$, either $c = H^{\text{T}}(\theta) = H(\theta)$ for all $\theta \in \Theta_c$, or $\int_{\Theta_c} [H^{\text{T}}(\theta) - H(\theta)] \, dF(\theta) = 0$. In either case, since $v(q^*(\theta))$ is constant on Θ_c , we have

$$\int_{\Theta_c} [H^{\text{T}}(\theta) - H(\theta)] v(q^*(\theta)) \, dF(\theta) = v(q^*(\theta)) \int_{\Theta_c} [H^{\text{T}}(\theta) - H(\theta)] \, dF(\theta) = 0,$$

so that integrating over all $c \in \text{im } H$ obtains the required result.

Second, we show that for any feasible q

$$\int_{\Theta} [H(\theta) - H^{\text{T}}(\theta)] v(q(\theta)) \, dF(\theta) \geq 0.$$

Intuitively, wherever $H \neq H^{\text{T}}$, H^{T} is nonmonotone, and so if q is increasing, it tends to be larger where H^{T} is smaller. Formally, we calculate the integral as

$$\left[v(q(\theta)) \int_{\underline{\theta}}^{\theta} [H(s) - H^{\text{T}}(s)] \, dF(s) \right]_{\theta=\underline{\theta}}^{\theta=\bar{\theta}} - \int_{\underline{\theta}}^{\bar{\theta}} \left\{ \frac{\partial}{\partial \theta} [v(q(\theta))] \int_{\underline{\theta}}^{\theta} [H(s) - H^{\text{T}}(s)] \, dF(s) \right\} \, d\theta,$$

which is nonnegative because $v(q(\theta))$ is increasing in θ for any feasible q and $\int_{\underline{\theta}}^{\theta} H^{\text{T}}(s) \, dF(s) \geq \int_{\underline{\theta}}^{\theta} H(s) \, dF(s)$ for all $\theta \in \Theta$ (with equality at $\underline{\theta}$ and $\bar{\theta}$) since the latter (viewed as a function of θ) is the convex envelope of the former.

Third, and finally, we show that for any feasible q

$$\int_{\Theta} [J(\theta) - H^{\text{T}}(\theta)] [v(q^*(\theta)) - v(q(\theta))] \, dF(\theta) \geq 0.$$

For this, we sign the two terms in the product. Because the target H^T is chosen to either equal J (in which case the first term is zero) or to equal the lower bound $L(\theta)$, we need only consider the integrand where the lower-bound constraint is binding, which is when $L(\theta) \geq \overline{J}_{[\underline{\theta}, \theta]}(\theta)$ and $H^T(\theta) = L(\theta)$. But by Lemma 6, $\overline{J}_{[\underline{\theta}, \theta]}(\theta) \geq J(\theta)$, so $L(\theta) \geq J(\theta)$, and thus $[J(\theta) - H^T(\theta)]$ is nonpositive. For the second term in the product, where $H^T(\theta) = L(\theta)$, we have via Lemma 7 that $H(\theta) = L(\theta)$, so $q^*(\theta) = q^L(\theta)$. As a result, $q^*(\theta) = q^L(\theta) \leq q(\theta)$ for any feasible q . By the monotonicity of v , we have that $v(q^*(\theta)) - v(q(\theta)) \leq 0$, so the integrand is nonnegative for all θ , completing the proof. \square

A.7 Proof of Proposition 4 and the Optimal Payment Schedule

Returning now to the case in which $L(\theta) = \theta$, we derive the payment schedule associated with the optimal mechanism. We first derive Proposition 4 which pins down the payment made by the lowest type $\underline{\theta}$.

Proof. As shown in Appendix A.2 above, when $\mathbf{E}[\omega] > \alpha$, the social planner benefits from making $U(\underline{\theta})$ as large as possible, which implies that the (NLS) constraint binds for $\underline{\theta}$ and the social planner offers the initial $q^*(\underline{\theta})$ units of the good for free.

On the other hand, when $\mathbf{E}[\omega] < \alpha$, the social planner benefits from making $U(\underline{\theta})$ as small as possible, which implies that the (IR $_{\theta}$) constraint binds for $\underline{\theta}$. Because $J(\underline{\theta}) < \theta$, we have $H(\underline{\theta}) = \underline{\theta}$, so $q^*(\underline{\theta}) = D(c, \underline{\theta}) = q^{\text{LF}}(\underline{\theta})$. The initial $q^*(\underline{\theta})$ units of the good are priced at c . More generally, the payment schedule as a function of total quantity can be determined as follows.

When $\mathbf{E}[\omega] = \alpha$, $U(\underline{\theta})$ does not enter the social planner's objective and $q^*(\underline{\theta}) = q^{\text{LF}}(\underline{\theta})$. As a result, the social planner is indifferent between feasible payments for $q^{\text{LF}}(\underline{\theta})$ units of the good, which by the (NLS) and (IR $_{\theta}$) constraints imply $0 \leq t(\underline{\theta}) \leq cq^{\text{LF}}(\underline{\theta})$. \square

Proposition 15 (payment schedule). *The following subsidized payment schedule implements the optimal subsidy allocation for quantities $z \in \text{im } q^*$:*

$$P^\sigma(z) = \begin{cases} \int_{q^*(\underline{\theta})}^z (\theta^*)^{-1}(\hat{z})v'(\hat{z}) \, d\hat{z} & \text{if } \mathbf{E}[\omega] > \alpha, \\ \gamma + \int_{q^{\text{LF}}(\underline{\theta})}^z (\theta^*)^{-1}(\hat{z})v'(\hat{z}) \, d\hat{z} & \text{if } \mathbf{E}[\omega] = \alpha, \\ cq^{\text{LF}}(\underline{\theta}) + \int_{q^{\text{LF}}(\underline{\theta})}^z (\theta^*)^{-1}(\hat{z})v'(\hat{z}) \, d\hat{z} & \text{if } \mathbf{E}[\omega] < \alpha, \end{cases}$$

where $(\theta^*)^{-1}(z) = \sup\{\theta \in \Theta : q^*(\theta) \geq z\}$ is the generalized inverse of the optimal allocation rule, and $\gamma \in [0, cq^{\text{LF}}(\underline{\theta})]$.

Outside of $\text{im } q^*$, the payment schedule is not uniquely determined: when $\mathbf{E}[\omega] < \alpha$, it suffices to price units below $q^*(\underline{\theta})$ and above $q^*(\bar{\theta})$ at the marginal price c ; whereas when $\mathbf{E}[\omega] \geq \alpha$, it suffices to price units below $q^*(\underline{\theta})$ at zero and above $q^*(\bar{\theta})$ at c .

Proof. By the [Milgrom and Segal \(2002\)](#) envelope theorem, we have

$$t^*(\theta) = \theta v(q^*(\theta)) - U(\underline{\theta}) - \int_{\underline{\theta}}^{\theta} v(q^*(s)) \, ds.$$

Because $q^*(\theta)$ is nondecreasing and bounded, it is differentiable almost everywhere, and so we have for almost all $z \in \text{im } q^* = [q^*(\underline{\theta}), q^*(\bar{\theta})]$ that

$$\frac{d}{dz} P^\sigma(z) = \frac{dt^*/d\theta}{dq^*/d\theta} = (\theta^*)^{-1}(z)v'(z),$$

which, on integrating and applying our previous observations about the payment made by the lowest type, gives the desired result. \square

Note that [Proposition 15](#) implies that the total subsidy rule satisfies

$$\begin{aligned} \sigma(z) &= c - \frac{d}{dz} P^\sigma(z) \\ &= c - (\theta^*)^{-1}(z)v'(z) \\ &= c - (\theta^*)^{-1}(z) \frac{c}{H((\theta^*)^{-1}(z))}, \end{aligned}$$

which is nonnegative because $H(\theta) \geq \theta$, as necessary for feasibility.

B Structure of Optimal Subsidies

In this section, we derive the structure of the optimal subsidy mechanism for increasing and decreasing ω , proving Proposition 5 and Proposition 6 stated in Section 5.2.

B.1 Proof of Proposition 5 (Negative Correlation)

To derive Proposition 5, we study the implications of a decreasing $\omega(\theta)$ on the virtual welfare $J(\theta)$, separately when the opportunity cost of funds is high and low, focusing on the sign of the distortion term, equal to $\frac{\int_{\underline{\theta}}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)}{\alpha f(\theta)}$ for $\theta \in (\underline{\theta}, \bar{\theta}]$.

High opportunity cost of funds: $\mathbf{E}[\omega] \leq \alpha$. When the opportunity cost of funds is high, Theorem 1 implies that the optimal mechanism is laissez-faire. To see this in terms of the construction of the subsidy type, note that a decreasing ω implies that $\int_{\underline{\theta}}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)$ is quasiconvex in θ , while it is zero at $\bar{\theta}$ and nonpositive at $\underline{\theta}$ when $\mathbf{E}[\omega] \leq \alpha$. As a result, the distortion term is nonpositive for all $\theta \in \Theta$, as illustrated in Figure 10.²²




Figure 10: Subsidy types for decreasing welfare weights and a high opportunity cost of funds. In this case, the lower-bound constraint binds for all types.

Low opportunity cost of funds: $\mathbf{E}[\omega] > \alpha$ Proposition 4 implies that there is always public provision of the good when $\mathbf{E}[\omega] > \alpha$. To determine the quantity of public provision and any

²² In Figures 10, 11, 12, and 13, we have depicted examples in which there are no non-monotonicities in $J(\theta)$ other than that caused by the point mass at $\underline{\theta}$. But even under monotonicity assumptions on $\omega(\theta)$, the virtual welfare $J(\theta)$ may contain additional non-monotonicities because the distortion term depends inversely on $f(\theta)$, which may also be non-monotone, possibly leading to additional ironing intervals in $H(\theta)$.

other subsidies, we construct the subsidy type, noting again that a decreasing ω implies that $\int_{\underline{\theta}}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)$ is quasiconvex in θ , while it is zero at $\bar{\theta}$ and positive at $\underline{\theta}$ when $\mathbf{E}[\omega] > \alpha$. This means there are two possibilities for the sign of the distortion term:


- (a) The distortion term is positive for all types $\theta \in \Theta$. This requires $\omega(\theta) \geq \alpha$ for all $\theta \in \Theta$, in which case, the social planner wants to distort the allocation of each type upward. The subsidy type equals the ironed virtual welfare, as illustrated in Figure 11(a), so $q^* = q^{\text{SD}}$.
- (b) There exists a type $\hat{\theta} \in \Theta$ such that the distortion term is nonnegative for $\theta \leq \hat{\theta}$ and nonpositive for all $\theta \geq \hat{\theta}$. In this case, the social planner wants to distort the allocation of types lower than $\hat{\theta}$ upward and distort the allocation of types greater than $\hat{\theta}$ downward. There are two possibilities either:

- (i) $\overline{J|_{[\underline{\theta}, \hat{\theta}]}}(\hat{\theta}) = J(\hat{\theta}) = \hat{\theta}$ (i.e., $\hat{\theta}$ is not in the interior of an ironing interval of J), in which case $H(\theta) = \overline{J}(\theta)$ on $[\underline{\theta}, \hat{\theta}]$ and then $H(\theta) = \theta$ on $[\hat{\theta}, \bar{\theta}]$, as shown in Figure 11(b)(i). Writing $\hat{\theta} = \theta_\alpha$ as in Table 1, $q^*(\theta) = q^{\text{SD}}(\theta)$ for $\theta \leq \theta_\alpha$, and $q^*(\theta) = q^{\text{LF}}(\theta)$ for $\theta \geq \theta_\alpha$, with the associated payment schedule derived using Proposition 15, incorporating quantity discounts between θ_1 (the end of the ironing interval of $J(\theta)$ starting at $\underline{\theta}$) and θ_α , as illustrated in Figure 1(a).
- (ii) $\overline{J|_{[\underline{\theta}, \hat{\theta}]}}(\hat{\theta}) > J(\hat{\theta}) = \hat{\theta}$ (i.e., $\hat{\theta}$ is in the interior of an ironing interval of J), in which case there is a least type $\theta_\alpha \geq \hat{\theta}$ for which $\overline{J|_{[\underline{\theta}, \theta_\alpha]}}(\theta_\alpha) = \theta_\alpha$. In that case, $H(\theta) = \overline{J|_{[\underline{\theta}, \theta_\alpha]}}(\theta)$ on $[\underline{\theta}, \theta_\alpha]$ and $H(\theta) = \theta$ on $[\theta_\alpha, \bar{\theta}]$, as in Figure 11(b)(ii). This implies $q^* \geq q^{\text{SD}}$, with strict inequality on the ironing interval of $J|_{[\underline{\theta}, \theta_\alpha]}$ ending at θ_α , and $q^* = q^{\text{LF}}$ on $[\theta_\alpha, \bar{\theta}]$. In this case, the payment schedule, derived using Proposition 15, involves a free allocation of $q^{\text{LF}}(\theta_\alpha)$ units for all agents and laissez-faire prices for additional consumption beyond that amount.

B.2 Proof of Proposition 6 (Positive Correlation)


We again derive Proposition 6 by examining the implications of increasing ω for the distortion term of the virtual welfare.

High opportunity cost of funds: $\mathbf{E}[\omega] \leq \alpha$. When ω is decreasing in θ , we have that $\int_{\underline{\theta}}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)$ is quasiconcave in θ , while it is zero at $\bar{\theta}$ and nonpositive at $\underline{\theta}$ because $\mathbf{E}[\omega] \leq \alpha$. This leads to two possibilities for the sign of the distortion term:




(a) Here $\omega(\theta) \geq \alpha$ for each θ , so the social planner distorts each type's allocation upward, with a free endowment for all types and quantity-dependent subsidies for additional units of the good.

— subsidy type $H(\theta)$ - - - virtual welfare $J(\theta)$ - - - lower bound θ



(b)(i) Here, the social planner distorts lower types' consumption upward, via a free endowment of $q^{\text{LF}}(J(\theta_1))$ and limited quantity-dependent subsidies for additional units. Consumption of types $\theta \geq \theta_\alpha$ is undistorted.




(b)(ii) In this example, the social planner distorts lower types' consumption upward, via a free endowment of $q^{\text{LF}}(\theta_\alpha)$, which can strictly exceed $q^{\text{SD}}(\theta)$ for $\theta \leq \theta_\alpha$. Consumption of types $\theta \geq \theta_\alpha$ is undistorted.


Figure 11: Subsidy types for decreasing welfare weights and low opportunity cost of funds.

- (a) The distortion term is nonpositive for all $\theta \in \Theta$, in which case $\max_{\theta} \omega(\theta) < \alpha$, and the social planner wants to distort all consumers' consumption downward, but is prevented from doing so by the (LB) constraint. As a result, $H(\theta) = \theta$, as illustrated in Figure 12(a), so $q^* = q^{\text{LF}}$, and the social planner offers no subsidies.
- (b) There exists a $\theta_\alpha \in \Theta$ such that the distortion term is nonpositive for $\theta \in [\underline{\theta}, \theta_\alpha]$ and nonnegative for $\theta \in [\theta_\alpha, \bar{\theta}]$. This means $H(\theta) = \bar{J}(\theta)$ for $\theta \geq \theta_\alpha$, and $H(\theta) = \theta$ for $\theta \leq \theta_\alpha$, as illustrated in Figure 12(b). As a result, $q^*(\theta) = q^{\text{LF}}(\theta)$ for $\theta \leq \theta_\alpha$, and $q^*(\theta) = q^{\text{SD}}(\theta)$ for $\theta \geq \theta_\alpha$. By Proposition 15, the social planner implements this allocation using subsidies for

consumption above a certain minimum level, as illustrated in Figure 1(b).



(a) In this case, the social planner wants to distort consumption for all types downwards, but is prevented from doing so by the (LB) constraint, so $q^* = q^{\text{LF}}$.




(b) In this case, the social planner wants to distort consumption for higher types upwards, which it does by offering discounts for consumption above a minimum level.

— subsidy type $H(\theta)$ - - - virtual welfare $J(\theta)$ - - - lower bound θ

Figure 12: Subsidy types with increasing welfare weights and a high cost of public funds.

Low opportunity cost of funds: $\mathbf{E}[\omega] > \alpha$. In this case, Proposition 4 implies that there is public provision of the good because $\mathbf{E}[\omega] > \alpha$. To determine the quantity of public provision and any other subsidies, we construct the subsidy type, noting again that an increasing ω implies that $\int_{\theta}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)$ is quasiconcave in θ , while it is zero at $\bar{\theta}$ and positive at $\underline{\theta}$ when $\mathbf{E}[\omega] > \alpha$. It is thus positive for all $\theta \in \Theta$, so the social planner wants to distort consumption of all types upward. As illustrated in Figure 13, the subsidy type equals the ironed virtual welfare, so $q^* = q^{\text{SD}}$.



— subsidy type $H(\theta)$ - - - virtual welfare $J(\theta)$ - - - lower bound θ

Figure 13: Subsidy types with increasing welfare weights and low opportunity cost of funds. Here, the social planner wants to distort consumption for all types upwards, so the social planner offers a free allocation with quantity-dependent subsidies for additional units of the good.

C Omitted proofs

C.1 Restriction to Deterministic Mechanisms

Suppose the social planner were to use a randomized mechanism, and let $M : \Theta \rightarrow \Delta([0, A] \times \mathbb{R})$ be the randomized allocation and payment rule, with $M(\theta)$ determining for a type θ a joint distribution over quantities and payments, assumed to be nondegenerate for some positive measure of types. The constraints facing the social planner are now:

- incentive-compatibility

$$\theta \in \arg \max_{\hat{\theta}} \mathbf{E}_{(q,t) \sim M(\hat{\theta})} [\theta v(q) - t],$$

- individual rationality,

$$\mathbf{E}_{(q,t) \sim M(\theta)} [\theta v(q) - t] \geq U^{\text{LF}}(\theta),$$

- no lump-sum transfers

$$\mathbf{E}_{(q,t) \sim M(\theta)} [t] \geq 0,$$

which we assume need only be satisfied in expectation over M (but the argument is more or less unchanged if all realizations of t must exceed 0), and

- the private market constraint

$$\Pr_{(q,t) \sim M(\theta)} [q \geq q^{\text{LF}}(\theta)] = 1,$$

which we enforce on all realizations of the lottery (assuming the consumer could always top up contingent on the outcome of the randomized mechanism).

We claim that the deterministic mechanism (\hat{q}, \hat{t}) in which $\hat{q}(\theta)$ is the certainty equivalent of the randomized quantity assigned in $M(\theta)$ and $\hat{t}(\theta)$ is the expected payment under $M(\theta)$ satisfies all the constraints and strictly improves the social planner's objective. To see this, note that by this construction, for all types $\theta, \hat{\theta} \in \Theta$, we have $\mathbf{E}_{(q,t) \sim M(\hat{\theta})} [\theta v(q) - t] = \theta v(\hat{q}(\hat{\theta})) - \hat{t}(\hat{\theta})$, which implies the usual incentive-compatibility constraint and individual rationality constraints are satisfied for (\hat{q}, \hat{t}) . The no-lump sum transfers constraint and private market constraints are also clearly satisfied.

To see that this construction results in a strict improvement in the social planner's objective, consider any type θ that received a nondegenerate lottery in M . Recall that for a strictly concave

valuation function, the certainty equivalent of any nondegenerate lottery $\hat{q}(\theta) < \mathbf{E}_{(q,t) \sim M(\theta)}[q]$. As a consequence, the weighted consumer surplus

$$\int_{\Theta} \omega(\theta) [\theta v(\hat{q}(\theta)) - \hat{t}(\theta)] + \alpha [\hat{t}(\theta) - c\hat{q}(\theta)] \, dF(\theta)$$

is strictly larger than the weighted consumer surplus under the randomized mechanism

$$\int_{\Theta} \omega(\theta) \mathbf{E}_{(q,t) \sim M(\theta)} [\theta v(q) - t] + \alpha \mathbf{E}_{(q,t) \sim M(\theta)} [t - cq] \, dF(\theta),$$

because the expected utility of each type θ is unchanged (as is his expected payment), and the expected cost for the social planner is strictly reduced (because fewer total goods are produced).

We note that the *strict* preference for deterministic mechanisms relies on the strict concavity of the consumer's valuation function. Without strict concavity, the social planner might be indifferent between the optimal mechanism identified in Theorem 2 and a mechanism including randomization.

C.2 Existence and Uniqueness of Optimal Mechanism

Existence. Note that the social planner's objective expressed in terms of q , (see (OPT- ℓ) for $\mathbf{E}[\omega] > \alpha$ and (OPT- h) for $\mathbf{E}[\omega] \leq \alpha$) is continuous in q (in the L^1 topology). By the Helly selection theorem, the set of nonincreasing and bounded functions $q : \mathbb{R} \rightarrow [0, A]$ is compact, so that the set of feasible allocation functions (a subset of that set) is compact as well. On the other hand, the set of feasible allocation rules contains q^{LF} and is thus nonempty. As a result, an optimal solution exists.

Uniqueness. We now show that the optimal solution to the social planner's problem is unique, assuming that consumers have a strictly concave valuation function. Suppose, for a contradiction, that q_1 and q_2 are distinct optimal solutions and consider \tilde{q} defined by $v \circ \tilde{q}(\theta) = \frac{1}{2} [v(q_1(\theta)) + v(q_2(\theta))]$. Clearly, \tilde{q} is nondecreasing and feasible because each constraint is linear in $v \circ q$. On the other hand, because v is strictly concave, the social planner's objective is a strictly convex function of $v \circ q$ (the integrand may be written $J(\theta)\nu(\theta) - cv^{-1}(\nu(\theta))$, where $\nu = v \circ q$ and v^{-1} is strictly convex). But then by Jensen's inequality, we have that the social planner's objective is strictly larger at \tilde{q} , contradicting the optimality of q_1 and q_2 .

C.3 Additional Details for Proof of Theorem 1 (When the Social Planner Uses Subsidies)

The key details of the proof of Theorem 1 are contained in Section 3, but we fill in two missing details: first, the general proof that there is no beneficial in-kind subsidy program for the planner when the condition in Theorem 1 is not satisfied, and second, we show that the total cost of the subsidy to consumers with types $\theta \leq \hat{\theta}$ is of lower order than the linear benefits to types $\theta > \hat{\theta}$.

For the first, suppose that for all types $\hat{\theta} \in \Theta$,

$$\mathbf{E}_\theta [\omega(\theta) | \theta \geq \hat{\theta}] \leq \alpha,$$

and suppose that the social planner pays strictly positive subsidies to some set of consumers. By Lemma 3, there is a least consumer type $\hat{\theta} \in \Theta$ receiving a nonnegative subsidy payment $\Sigma(q(\hat{\theta}))$, and all types $\theta \geq \hat{\theta}$ must receive at least as large a subsidy payment.

Consider

$$\int_{\Theta} \omega(\theta)[U(\theta) - U^{\text{LF}}(\theta)] - \alpha \Sigma(q(\theta)) \, dF(\theta),$$

which is an expression for the weighted total surplus associated with marginal subsidy rule σ , minus the weighted consumer surplus under the laissez-faire allocation rule. Note that this is the social planner's objective minus a term independent of the subsidy mechanism chosen, so it is equivalent to formulate the problem in terms of a choice of subsidy program to maximize this expression. On the other hand, this objective is zero for the laissez-faire mechanism, so to show that the social planner does not benefit from subsidies, it suffices to show that the objective above is bounded above by zero.

Because $U(\theta) - U^{\text{LF}}(\theta) \leq \Sigma(q(\theta))$ (as shown in Section 3.1), we have that this expression is bounded by

$$\int_{\Theta} [\omega(\theta) - \alpha] \Sigma(q(\theta)) \, dF(\theta).$$

Integrating by parts, we may rewrite this as

$$\Sigma(q(\hat{\theta})) \int_{\hat{\theta}}^{\bar{\theta}} [\omega(s) - \alpha] \, dF(s) + \int_{\hat{\theta}}^{\bar{\theta}} \int_{\theta}^{\bar{\theta}} [\omega(s) - \alpha] \, dF(s) \, d\Sigma(q(\theta)),$$

which is nonpositive because $\Sigma \circ q$ is increasing (by Lemma 3) and $\int_{\theta}^{\bar{\theta}} \omega(s) - \alpha \, dF(s)$ is nonpositive by assumption.

We now fill in the second remaining detail, showing that the total cost of the subsidy to

consumers with types $\theta \leq \hat{\theta}$ as defined in Section 3.1 is of lower order than the linear benefits to types $\theta > \hat{\theta}$.

To see this, we calculate the type $\tilde{\theta}$ just indifferent to distorting its consumption to receive the subsidy, which means

$$\tilde{\theta}v(q^{\text{LF}}(\hat{\theta})) - cq^{\text{LF}}(\hat{\theta}) + \varepsilon = \tilde{\theta}v(q^{\text{LF}}(\tilde{\theta})) - cq^{\text{LF}}(\tilde{\theta}).$$

By monotonicity of demand, the set of types with distorted consumption under the ε -perturbed subsidy schedule is $(\tilde{\theta}, \hat{\theta})$. We thus bound $\hat{\theta} - \tilde{\theta}$.

Because $v'' < 0$, we have that there exists a $k > 0$ such that $v'' < -k$ on $\text{int } \Theta$. This implies that v is strongly concave (cf. [Watt \(2022\)](#)), so

$$\tilde{\theta}v(q^{\text{LF}}(\hat{\theta})) \leq \tilde{\theta}v(q^{\text{LF}}(\tilde{\theta})) + c(q^{\text{LF}}(\hat{\theta}) - q^{\text{LF}}(\tilde{\theta})) - k[q^{\text{LF}}(\hat{\theta}) - q^{\text{LF}}(\tilde{\theta})]^2,$$

which implies

$$k[q^{\text{LF}}(\hat{\theta}) - q^{\text{LF}}(\tilde{\theta})]^2 \leq \varepsilon.$$

On the other hand,

$$\frac{\partial}{\partial \theta} D(c, \theta) = \frac{-c}{\theta^2 v''[(v')^{-1}(c/\theta)]} \geq \frac{c}{k\underline{\theta}^2},$$

so

$$q^{\text{LF}}(\hat{\theta}) - q^{\text{LF}}(\tilde{\theta}) \geq \frac{c}{k\underline{\theta}^2}[\hat{\theta} - \tilde{\theta}].$$

Putting these inequalities together, we obtain

$$\hat{\theta} - \tilde{\theta} \leq \sqrt{\frac{\varepsilon}{k} \frac{k\underline{\theta}^2}{c}} \sim O(\sqrt{\varepsilon}).$$

But then since F is absolutely continuous, we have that $F(\hat{\theta}) - F(\tilde{\theta}) \sim O(\sqrt{\varepsilon})$ so the total costs (bounded below by $-\alpha\varepsilon$ per consumer) are $O(\varepsilon^{\frac{3}{2}})$. This means that the benefits are linear in ε while the costs are $O(\varepsilon^{\frac{3}{2}})$, which means the net effect must be positive for sufficiently small ε .

C.4 Shutdown Benchmark

In this section, we derive the solution to the *shutdown benchmark* problem, in which the social planner can foreclose the private market for the good.

In that case, the social planner's problem takes the form:

$$\max_{(q,t)} \int_{\Theta} \left\{ \omega(\theta) \underbrace{[\theta v(q(\theta)) - t(\theta)]}_{\text{consumer surplus}} + \alpha \underbrace{[t(\theta) - cq(\theta)]}_{\text{total profit}} \right\} dF(\theta), \quad (\text{NO-PM})$$

such that (q, t) satisfies (IC), (NLS), and for all $\theta \in \Theta$, $U(\theta) \geq 0$.

This can be rewritten as in Appendix A.2 as

$$\max_{q \in \mathcal{Q}} \alpha \int_{\Theta} [J(\theta)v(q(\theta)) - cq(\theta)] dF(\theta),$$

which differs from (OPT- q) only in the lack of a lower-bound constraint and the absence of the terms outside the integrand (independent of q) arising from the stronger individual rationality constraints in the subsidy design problem (that is, individual rationality requires only that each consumer's utility exceeds zero, rather than their laissez-faire utility).

We now show that the optimal allocation rule in this problem is

$$q^{\text{SD}}(\theta) = (v')^{-1} \left(\frac{c}{\bar{J}(\theta)} \right).$$

To see this, note that q^{SD} is the pointwise maximizer of the integrand

$$\int_{\Theta} [\bar{J}(\theta)v(q(\theta)) - cq(\theta)] dF(\theta),$$

so, via similar logic to Appendix A.6, it suffices to show that for any feasible $q \in \mathcal{Q}$,

$$\int_{\Theta} [J(\theta) - \bar{J}(\theta)][v(q^{\text{SD}}(\theta)) - v(q(\theta))] dF(\theta) \geq 0.$$

But $J(\theta) = \bar{J}(\theta)$ except on ironing intervals of J , in which q^{SD} is constant and \bar{J} is the F -average of J , so

$$\int_{\Theta} [J(\theta) - \bar{J}(\theta)] v(q^{\text{SD}}(\theta)) dF(\theta) = 0,$$

as we showed in Appendix A.6. Finally, as we also showed in Appendix A.6, because any feasible q is nondecreasing wherever $J \neq \bar{J}$, we have that

$$\int_{\Theta} [J(\theta) - \bar{J}(\theta)] v(q^{\text{SD}}(\theta)) dF(\theta) \leq 0,$$

completing the proof.

C.5 Linear Subsidies

Optimal Linear Subsidies and Proposition 3. In this section, we establish Proposition 3 and its generalizations without Assumption 1 and 2.

Proof. By the same logic as that presented in Appendix A.2, we can write the problem of a social planner restricted to linear subsidies as

$$\max_{s \in [0,1]} \left\{ (\mathbf{E}[\omega] - \alpha) \underline{U}(s) + \alpha \int_{\Theta} \left[\left(\theta + \frac{\int_{\theta}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)}{\alpha f(\theta)} \right) v(D(cs, \theta)) - cD(cs, \theta) \right] dF(\theta) \right\}$$

Write the coefficient outside of v as $\tilde{J}(\theta)$, in analogy to the virtual type. By the envelope theorem (or Roy's Identity), $\underline{U}'(s) = -cD(cs, \theta)$, so the derivative of the objective in s is

$$-(\mathbf{E}[\omega] - \alpha)cD(cs, \theta) + \alpha \int_{\Theta} \left[\tilde{J}(\theta)v'(D(cs, \theta)) \frac{dD(cs, \theta)}{ds} - c \frac{dD(cs, \theta)}{ds} \right] dF(\theta),$$

which on calculating the derivative of demand can be rewritten

$$-(\mathbf{E}[\omega] - \alpha)cD(cs, \theta) + \alpha c^2 \int_{\Theta} \frac{s\tilde{J}(\theta) - \theta}{\theta^2 v'' \circ (v')^{-1}(cs/\theta)} dF(\theta).$$

The optimal subsidy thus satisfies

$$-(\mathbf{E}[\omega] - \alpha)cD(cs, \theta) + \alpha c^2 \int_{\Theta} \frac{s\tilde{J}(\theta) - \theta}{\theta^2 v'' \circ (v')^{-1}(cs/\theta)} dF(\theta) = 0.$$

If the social planner's problem is concave in s , the condition for intervention (i.e., $s^* < 1$) is then

$$-(\mathbf{E}[\omega] - \alpha)cD(c, \theta) + \alpha c^2 \int_{\Theta} \frac{\tilde{J}(\theta) - \theta}{\theta^2 v'' \circ (v')^{-1}(c/\theta)} dF(\theta) \leq 0.$$

To obtain Proposition 3, we note that under Assumption 1 and 2, the general expression for the derivative of the objective is decreasing in s (so the social planner's problem is concave), the first term is zero, and $v'' = -1$, leading to the expression in Proposition 3. \square

Optimal Nonlinear Subsidies Versus Optimal Linear Subsidies: Proposition 7 and 8.

We now prove Proposition 7, establishing that the optimal nonlinear subsidy mechanism is never linear.

Proof. Recall that $J(\bar{\theta}) = \bar{\theta}$. This implies that $H(\bar{\theta})$ is either equal to $\bar{\theta}$, so the consumption of that type is undistorted in the optimal subsidy mechanism, or $H(\bar{\theta}) > \bar{\theta}$ and equal to the average value of J over some interval $[\hat{\theta}, \bar{\theta}]$. The first possibility can never be achieved using linear subsidies, which distorts *all* types' consumption upwards. The second possibility implies that consumption is constant for some set of types $[\hat{\theta}, \bar{\theta}]$, which also never occurs in a linear subsidy mechanism because the objective $\theta v(q) - csq$ is strictly supermodular (cf. Edlin and Shannon (1998)) and so $q(\theta)$ is strictly increasing. \square

We now establish Proposition 8, showing that it is always possible to identify a simple improvement over a linear subsidy mechanism.

Proof. For this proof, write $q^S(\theta) = D(cs, \theta)$.

Let us first consider the positive correlation case. As discussed in the proof of Proposition 5, the distortion term in this case is either always negative, negative then positive or always positive, depending on the sign of $\mathbf{E}[\omega] - \alpha$.

If the distortion term is always negative (i.e., if $\omega(\bar{\theta}) \leq \alpha$), then Theorem 1 implies the social planner is best off not offering subsidies at all, which can be achieved by setting a very high floor (or copay) on the subsidy, resulting in no consumers being subsidized.

If $\mathbf{E}[\omega] < \alpha$ but $\omega(\bar{\theta}) > \alpha$, then there exists a largest type θ_α such that $J(\theta) \leq \theta_\alpha$ for all $\theta \leq \theta_\alpha$. In that case, by setting a copay on the subsidy, so $\sigma'(q) = 0$ for all $q \leq q^S(\hat{\theta})$, all types below θ_α have their consumption reduced to some $\hat{q}(\theta)$ with $q^{\text{LF}}(\theta) \leq \hat{q}(\theta) \leq q^S(\theta)$, while all types above θ_α continue to consume $q^S(\theta)$. In that case, the increase in the objective value caused by the change is

$$\int_{\underline{\theta}}^{\theta_\alpha} J(\theta)[v(\hat{q}(\theta)) - v(q^S(\theta))] - c[\hat{q}(\theta) - q(\theta)] dF(\theta),$$

which, given $J(\theta) \leq \theta$ and $q^{\text{LF}}(\theta) \leq \hat{q}(\theta) \leq q^S(\theta)$, is at least

$$\int_{\underline{\theta}}^{\theta_\alpha} \theta[v(\hat{q}(\theta)) - v(q^S(\theta))] - c[\hat{q}(\theta) - q^S(\theta)] dF(\theta) \geq 0,$$

by definition of q^{LF} . We have thus identified an improvement in the social planner's objective.

If $\mathbf{E}[\omega] = \alpha$, then $J(\underline{\theta}) = \underline{\theta}$ and for θ in a neighborhood of $\underline{\theta}$, $J(\theta)$ is in a neighborhood of $\underline{\theta}$. This means that the social planner would like to design a subsidy that leaves $\underline{\theta}$'s consumption undistorted and leads to arbitrarily small distortions for consumers in a neighborhood of $\underline{\theta}$. To achieve that goal, for any given s , the social planner can identify a copay such that the consumer of type $\underline{\theta}$ is just indifferent between consuming $q^{\text{LF}}(\underline{\theta})$ units at price c and consuming $q^S(\underline{\theta})$ units with the copay. In that case, by setting a copay ε larger for sufficiently small $\varepsilon > 0$, the social planner can distort consumption of types in a neighborhood of $\underline{\theta}$ down to θ , while leaving consumption outside of that neighborhood equal to $q^S(\underline{\theta})$. By similar logic to the above, this leads to an increase in the social planner's objective.

On the other hand, if $\mathbf{E}[\omega] > \alpha$, then the social planner benefits from offering the first $q^S(\underline{\theta})$ units for free and then continuing to subsidize the additional units, which does not change the allocation rule, but does increase the utility of the worst-off type, improving the social planner's objective.

Now, consider the negative correlation case. As discussed in the proof of Proposition 6, the distortion term in this case is either always negative (in which case the social planner should set a low cap on the subsidy resulting in no consumer subsidies), positive then negative or always positive (in which case our argument in the previous section for offering a free allocation continues to apply).

Thus, let us consider the only remaining case which is when the distortion term changes sign from positive to negative at some θ_α in the interior of Θ . In that case, by setting $\sigma'(q) = 0$ for $q \geq q^S(\theta_\alpha)$, consumers of types $\theta \geq \theta_\alpha$ consume at some $\hat{q}(\theta)$ with $q^{\text{LF}}(\theta) \leq \hat{q}(\theta) \leq q^S(\theta)$, while consumers of types $\theta < \theta_\alpha$ continue to consume $q^S(\theta)$ units of the good. By analogous logic to the one presented above for the negative to positive sign change, this constitutes an improvement in the social planner's objective.

□

C.6 Comparative Statics Proofs

To prove the comparative statics stated in Section 5.5, we characterize the change in virtual welfare caused by the change in the economic primitive and examine the effect of that change on the subsidy type and the resulting optimal subsidy mechanism. Our analyses use the following lemma, establishing the monotonicity of the ironing operator with respect to pointwise increases in its argument.

Lemma 8 (monotonicity of ironing operator). *Let J, J' be real-valued functions on Θ with $J'(\theta) \geq J(\theta)$ for all $\theta \in \Theta$. Then $\overline{J}'(\theta) \geq \overline{J}(\theta)$ for all $\theta \in \Theta$.*

Proof. Recall (see, e.g., [Reid \(1968\)](#), [Barron \(1983\)](#), and [Kang \(2024\)](#)) that \overline{J} solves

$$\max_{q \in \mathcal{Q}} \int_{\Theta} -[J(\theta) - q(\theta)]^2 dF(\theta).$$

Consider the parametrized objective

$$\max_{q \in \mathcal{Q}} \int_{\Theta} -[\alpha J'(\theta) + (1 - \alpha)J(\theta) - q(\theta)]^2 dF(\theta),$$

for $\alpha \in [0, 1]$ and write the solution of that problem for a fixed α as q_{α}^* . By the above, we have that $q_1^* = \overline{J}'$ and $q_0^* = \overline{J}$.

Let \succeq be the pointwise dominance partial order on the lattice \mathcal{Q} , so that $q \succeq q'$ if and only if $q(\theta) \geq q'(\theta)$ for all $\theta \in \Theta$. We claim that q_{α}^* is increasing in α in the pointwise dominance partial order, which implies that $\overline{J}'(\theta) \geq \overline{J}(\theta)$ for all $\theta \in \Theta$.

To establish this, note that the derivative of the parametrized objective with respect to α is

$$\int_{\Theta} -2[\alpha J'(\theta) + (1 - \alpha)J(\theta) - q(\theta)][J'(\theta) - J(\theta)] dF(\theta),$$

which is increasing in q in the pointwise dominance partial order. As a result, the objective is supermodular in (q, α) , so that q_{α}^* is increasing in the pointwise dominance partial order by the Topkis Theorem (cf. [Milgrom and Shannon, 1994](#)). \square

C.6.1 Proof of Proposition 9 (Increasing Redistributive Preferences)

Proof. The increase in $\omega(\theta)$ (in the sense of pointwise dominance) or the reduction in α leads to a pointwise increase in the distortion term $\int_{\theta}^{\overline{\theta}} [\omega(s) - \alpha] dF(s)$ in the virtual welfare and thus a pointwise increase in J . As a result, the value of the social planner's objective increases for every allocation rule $q(\cdot)$, and thus the value of the optimal mechanism also increases.

By Lemma 8, we have that $\overline{J}|_{\theta, \theta}(\theta)$ increases as a function of θ . As a result, the set of types with $\overline{J}|_{\theta, \theta}(\theta) > \theta$ (which is exactly the set of subsidized types) increases in the sense of inclusion, and the target H^T (introduced in Appendix A) also increases pointwise.

Applying Lemma 8 again, the increase in the target H^T implies that the subsidy type $H(\theta)$ increases for each type θ . Because each type's allocation is an increasing function of $H(\theta)$, that

implies that each agent's allocation increases as well. By the envelope theorem, the consumer surplus of each type θ also increases.

Finally, we show that the subsidy received by each type increases. To do so, we show an even stronger result: that the subsidy schedule, $\Sigma(q)$, becomes more generous. Because we established above that the allocation of each type also increases, that implies that the subsidy received by each type also increases. We argue in two steps: first, we show that the subsidy received by the lowest type increases (while the consumption of that type increases), and second, we show that the subsidy increases more quickly as a function of q .

For the lowest type, if $\mathbf{E}[\omega] \leq \alpha$ before the change, then it received no subsidy, and so any pointwise increase in the distortion term can only increase the subsidy received on the first $q^*(\underline{\theta})$ units (in particular, if $\mathbf{E}[\omega] > \alpha$ after the change, otherwise the consumption level is also unchanged). If $\mathbf{E}[\omega] > \alpha$ before the change, then the pointwise increase in the distortion term leads to an increase in his *free* allocation, which is an increase in the subsidy received for the first $q^*(\underline{\theta})$ units.

To see that the subsidized payment schedule is steeper, note that by Proposition 15

$$\frac{d}{dz}\Sigma(z) = c - (\theta^*)^{-1}(z)v'(z),$$

and $(\theta^*)^{-1}(z)$ is reduced pointwise (because $q(\theta)$ is increasing and increases pointwise). \square

C.6.2 Proof of Proposition 10 (Change in Interdependence)

Proof. As discussed in Section 5.5, the increase in the sense of majorization leads to a pointwise decrease in the virtual welfare function. As a result, the reverse of the analysis in the proof of Proposition 9 applies, leading to a less generous subsidy mechanism. \square

C.6.3 Proof of Proposition 11 (Increasing Demand)

Proof. A change in demand or costs results in no change to the virtual welfare weights, and so by Theorem 2, there is no change in the subsidy type. This implies there is no change in the set of types subsidized.

We now show that a pointwise increase in $v' : [0, A] \rightarrow \mathbb{R}_+$ or a decrease in c leads to an increase in $q^*(\theta) = (v')^{-1}(c/H(\theta))$. For the decrease in c , it suffices to note that

$$\frac{\partial}{\partial c}(v')^{-1}\left(\frac{c}{H(\theta)}\right) = \frac{1}{H(\theta)v''((v')^{-1}(c/H(\theta)))}$$

which is nonpositive by the strict concavity of v . For the pointwise increase in the marginal valuation v' , note that since v' is a nonincreasing function (because v is strictly concave), a pointwise increase in v' leads to a pointwise increase in $(v')^{-1}$, which implies a pointwise increase in q^* . By the envelope theorem, this also implies an increase in the consumer surplus of each type. \square

C.7 Proofs of Extensions

C.7.1 Proof of Proposition 12 (Equilibrium Effects)

Proof. As in the case with the constant marginal cost c , we can rewrite the objective eliminating the (IC) constraint as follows:

$$\max_{q \in \mathcal{Q}} \mathbf{E}[\omega - \alpha]U(\underline{\theta}) + \int_{\Theta} \left(\theta - \frac{\int_{\theta}^{\bar{\theta}} \omega(s) - \alpha \, dF(s)}{\alpha f(\theta)} \right) v(q(\theta)) \, dF(\theta) - \alpha C \left(\int_{\Theta} q(\theta) \, dF(\theta) \right),$$

subject to the (IR'), (NLS), and (PM'). Given any allocation function q and the price $p = S^{-1} \left(\int_{\Theta} q(\theta) \, dF(\theta) \right)$ it effects in the private market, the (IR') and (NLS) constraints can be eliminated and the objective can be rewritten as

$$\max_{q \in \mathcal{Q}} \int_{\Theta} J(\theta)v(q(\theta)) - pq(\theta) \, dF(\theta) + \int_{\Theta} pq(\theta) \, dF(\theta) - \alpha C \left(\int_{\Theta} q(\theta) \, dF(\theta) \right) + (\text{terms independent of } q),$$

for the same $J(\theta)$ defined as in the main analysis, subject to the constraint that $q(\theta) \geq (v')^{-1}(p/\theta) = D(p, \theta)$.

Replacing $J(\theta)$ with the subsidy type as defined in Theorem 2, we obtain the related problem

$$\max_{q \in \mathcal{Q}} \int_{\Theta} H(\theta)v(q(\theta)) - pq(\theta) \, dF(\theta) + \int_{\Theta} pq(\theta) \, dF(\theta) - \alpha C \left(\int_{\Theta} q(\theta) \, dF(\theta) \right),$$

subject to $q(\theta) \geq D(p, \theta)$ and $p = S^{-1} \left(\int_{\Theta} q(\theta) \, dF(\theta) \right)$. But by the First Welfare Theorem, we have that q^* defined as in Proposition 12 is the maximizer of this objective (which is the Paretian welfare objective for consumers with preferences determined by their subsidy type $H(\theta)$).

It thus remains to show that the maximizer of this expression is the same as the maximizer of the actual objective. But, as in the proof of Theorem 2, this involves showing for all feasible q that

$$\begin{aligned}
& \int_{\Theta} [J(\theta)v(q^*(\theta)) - pq^*(\theta)] dF(\theta) + \int_{\Theta} pq^*(\theta) dF(\theta) - \alpha C \left(\int_{\Theta} q^*(\theta) dF(\theta) \right) \\
& \geq \\
& \int_{\Theta} [J(\theta)v(q(\theta)) - pq(\theta)] dF(\theta) + \int_{\Theta} pq(\theta) dF(\theta) - \alpha C \left(\int_{\Theta} q(\theta) dF(\theta) \right),
\end{aligned}$$

given our previous finding that

$$\begin{aligned}
& \int_{\Theta} [H(\theta)v(q^*(\theta)) - pq^*(\theta)] dF(\theta) + \int_{\Theta} pq^*(\theta) dF(\theta) - \alpha C \left(\int_{\Theta} q^*(\theta) dF(\theta) \right) \\
& \geq \\
& \int_{\Theta} [H(\theta)v(q(\theta)) - pq(\theta)] dF(\theta) + \int_{\Theta} pq(\theta) dF(\theta) - \alpha C \left(\int_{\Theta} q(\theta) dF(\theta) \right),
\end{aligned}$$

which requires only that

$$\int_{\Theta} [J(\theta) - H(\theta)] [v(q^*(\theta)) - v(q(\theta))] dF(\theta) \geq 0,$$

which is exactly what we have shown in the proof of Theorem 2. \square

C.7.2 Proof of Proposition 13 (Exogenous Taxation)

Proof. By the linearity of demand, the lower-bound constraint (LB_t) can be written $q(\theta) \geq D(c, \frac{\theta}{1+\tau})$. The optimal allocation rule then follows directly from Theorem 2', replacing $L(\theta)$ with $\frac{\theta}{1+\tau}$. \square

D Additional Material

D.1 Optimal Subsidies in Pictures: A Two-Type Example

In this appendix, we present a two-type example to clarify the key tradeoffs in subsidy design and to build intuition for the optimal mechanism. We answer three key questions facing the social planner: (i) how can the social planner increase consumer surplus most cost-effectively using in-kind subsidies; (ii) when are such subsidies optimal; and (iii) how generous should they be?

For simplicity, we assume only in this section that there are low-demand and high-demand consumers, with types θ_L and θ_H respectively (where $\theta_L < \theta_H$); we refer to them as L -type and H -type consumers. The social planner assigns welfare weights ω_L and ω_H to these consumer types, with the proportion of L -type consumers being π_L , resulting in an average welfare weight $\mathbf{E}[\omega] = \pi_L \omega_L + (1 - \pi_L) \omega_H$.

Without subsidies, L -type consumers purchase $q_L^{\text{LF}} = D(c, \theta_L)$ units and H -type consumers purchase $q_H^{\text{LF}} = D(c, \theta_H)$ units of the good, as illustrated in Figure 14.




Figure 14: Illustrating the subsidy design problem using indifference curves.

The social planner chooses price-quantity pairs (p_L, q_L) and (p_H, q_H) for the L - and H -type consumers, with the subsidy paid equal to the vertical distance between the laissez-faire payment

schedule and the chosen point. Individual rationality requires that (p_L, q_L) lies on or below \tilde{U}_L^{LF} and that (p_H, q_H) lies on or below \tilde{U}_H^{LF} , and the (NLS) constraint requires p_L and p_H to be nonnegative. Fixing a choice of (p_L, q_L) , the H -type consumers' ability to access the private market price for the good requires (p_H, q_H) to lie in the gray shaded region in Figure 14 or on its boundary.

How Does The Social Planner Deliver Subsidies?

We first study how the social planner can most cost-effectively increase a consumer's surplus through subsidies.

The least costly way for the social planner to increase a consumer's surplus is through a *nondistortive cash transfer*, meaning a cash payment tied to the laissez-faire consumption of the good. This can be seen in Figure 14: the social planner minimizes the subsidy (the vertical distance between the laissez-faire payment schedule and the price-quantity pair) on \tilde{U}_L' , for example, at q_L^{LF} (as a result of the tangency condition for q_L^{LF} on \tilde{U}_L^{LF}).

However, the social planner's ability to offer nondistortive cash transfers is limited by the constraints in the subsidy design problem. For L -type consumers, the social planner can only provide a nondistortive cash transfer of cq_L^{LF} before violating the (NLS) constraint. Moreover, each dollar paid to L -type consumers lowers the blue line in Figure 14, requiring the social planner to reduce the price paid by the H -type by the same amount. For the H -type, the scope of nondistortive cash transfers depends on the allocation chosen for L -type consumers: the social planner can only offer nondistortive cash subsidies to H -type consumers when there is room between L -type consumers' indifference curve given their subsidy level, \tilde{U}_L' , and the blue line at q_H^{LF} .

To offer more generous subsidies while satisfying these constraints, the social planner must distort the consumption of one or both types of consumers upwards. Each dollar of consumer surplus gained via such *distortive subsidies* costs the social planner more than a dollar, as consumers value a dollar's worth of consumption beyond the laissez-faire amount less than a dollar. Moreover, due to the strict concavity of the consumer's preferences, that cost increases in the consumption distortion.²³ As a result, the social planner chooses the least distortive allocation for its desired subsidy spending while satisfying the problem's constraints. For the L -type, the binding constraint is the (NLS) constraint. For the H -type, the binding constraint is the (IC) constraint, reflecting the L -type's incentive to deviate upwards.

²³ Another implication of this fact is that the marginal benefit to the social planner of additional units of subsidy spending is decreasing.

When Does The Social Planner Offer Subsidies At All?

We now study when the social planner benefits from subsidizing consumption, first determining when the social planner subsidizes the L -type consumer. While the direct cost of subsidizing the L -type to choose (p_L, q_L) is the vertical distance between that point and the laissez-faire payment schedule, there is also an indirect cost: subsidizing L -type consumers lowers the blue line in Figure 14, making it possible for the H -type to imitate the L -type and purchase additional units in the private market, forcing the social planner to offer at least the same subsidy to the H -type. This means the social planner must trade off the benefit of increasing both types' ω -weighted consumer surplus against the α -weighted cost of the subsidy. As discussed above, the social planner can offer a small subsidy to L -type consumers without violating the (NLS) constraint or distorting consumption, resulting in an equal increase in consumer surplus for both types. Therefore, the social planner benefits from the initial units of subsidy for the L -type if and only if $\mathbf{E}[\omega] > \alpha$.

On the other hand, if $\mathbf{E}[\omega] < \alpha$, the social planner may still find it beneficial to subsidize the H -type but not the L -type. By tying a small subsidy to the H -type's consumption of q_H^{LF} units of the good, the social planner can increase the H -type's surplus without distorting consumption or violating the L -type's (IC) constraint. Thus, the initial units of subsidy for the H -type benefit the social planner if and only if $\omega_H > \alpha$.

In summary, the social planner offers subsidies if and only if $\mathbf{E}[\omega] > \alpha$ or $\omega_H > \alpha$.

How Generous Should Subsidies Be?

We now determine the optimal subsidy the social planner offers each consumer. We focus on two cases that illustrate the main tradeoffs facing the social planner:

- (a) *Decreasing welfare weights*, with $\omega_L > \alpha > \omega_H$ and $\mathbf{E}[\omega] > \alpha$. In this case, the generosity of subsidies offered to L -type consumers is constrained by both the (NLS) constraint and the downward incentive constraint of H -type consumers, who can mimic the L -type and then top up in the private market; and
- (b) *Increasing welfare weights*, with $\omega_L < \alpha < \omega_H$ and $\mathbf{E}[\omega] \leq \alpha$. In this case, the generosity of subsidies offered to H -type consumers is constrained by the upward incentive constraint of L -type consumers, who may prefer to distort their consumption upwards to receive a subsidy.

We then turn to the analysis of the only remaining case in which the social planner offers subsidies,

which is when $\mathbf{E}[\omega] > \alpha$ and $\omega_H > \alpha$, which combines features of both cases—namely, the binding (NLS) constraint from (a) and the binding upward (IC) constraint from (b).

Decreasing Welfare Weights: $\omega_L > \alpha > \omega_H$ with $\mathbf{E}[\omega] > \alpha$ In this case, because $\mathbf{E}[\omega] > \alpha$, every dollar of nondistortive subsidy spending for the L -type strictly increases the social planner's objective. Therefore, the optimal subsidy for the L -type is at least cq_L , leading to an optimal price-quantity pair of the form $(0, q_L)$ for some $q_L \geq q_L^{\text{LF}}$, to be determined.

However, since $\omega_H < \alpha$, the social planner does not benefit from further subsidizing the H -type consumer's consumption. This means, as illustrated in Figure 15(a), if $q_L \leq q_H^{\text{LF}}$, the downward (IC) constraint binds for the H -type and the social planner assigns the H -type $(c(q_H^{\text{LF}} - q_L), q_H^{\text{LF}})$. If $q_L > q_H^{\text{LF}}$ as in Figure 15(b), the (NLS) constraint binds for the H -type, and the social planner assigns $(0, q_L)$ to the H -type as well.

To determine q_L , the social planner must consider how her choice affects the constraints of the H -type. She evaluates the marginal benefit of each additional unit of subsidized consumption for the L -type as follows:

- (i) For $q_L \leq q_H^{\text{LF}}$, the marginal benefit is $\pi_L \omega_L \theta_L v'(q_L) + \pi_H \omega_H c - \alpha c$.
- (ii) For $q_L > q_H^{\text{LF}}$, the marginal benefit is $\mathbf{E}[\omega \theta] v'(q_L) - \alpha c$.

The optimal subsidy occurs when the marginal benefit of an additional unit of subsidized consumption is zero, leading to two possibilities either:

- (i) if $\pi_L \omega_L \theta_L v'(q_H^{\text{LF}}) + \pi_H \omega_H c < \alpha c$, then

$$q_L^* = (v')^{-1} \left(\frac{(\alpha - \pi_H \omega_H) c}{\pi_L \omega_L \theta_L} \right) \text{ and } p_L^* = 0,$$

and $q_H^* = q_H^{\text{LF}}$ and $p_H^* = c(q_H^* - q_L^*)$, as illustrated in Figure 15(a), or

- (ii) if $\pi_L \omega_L \theta_L v'(q_H) + \pi_H \omega_H c \geq \alpha c$, then there is pooling with

$$q_H^* = q_L^* = (v')^{-1} \left(\frac{\alpha c}{\mathbf{E}[\omega \theta]} \right),$$

and both types receive that quantity for free, as illustrated in Figure 15(b).

Holding the other parameters fixed, case (i) occurs for lower values of ω_L while case (ii) occurs when ω_L is large enough. This means the social planner pools consumption only when her preference to redistribute to the L -type is strong enough.




Figure 15: Illustrating the binding constraints when $\mathbf{E}[\omega] > \alpha$ and $\omega_H < \alpha$.

Summing up, the optimal subsidy mechanism with decreasing welfare weights subsidizes consumption up to some $q_L^* > q_L^{LF}$ and allows H -type consumers to top up in the private market, if desired.

Increasing Welfare Weights: $\omega_L < \alpha < \omega_H$ with $\mathbf{E}[\omega] < \alpha$ In this case, because $\mathbf{E}[\omega] < \alpha$, even nondistortive spending for the L -type strictly reduces the social planner's objective. Therefore, the social planner does not subsidize the L -type, meaning $q_L^* = q_L^{LF}$ and $p_L^* = cq_L^{LF}$.

The social planner can offer nondistortive cash subsidies to the H -type only until the L -type's upward (IC) constraint binds, which happens when $\theta_L v(q_H^{LF}) - p_H = \theta_L v(q_L^{LF}) - cq_L$. Since $\omega_H > \alpha$, each dollar spent on the H -type without distorting consumption benefits the social planner. However, the social planner can do better by using distortive subsidies, chosen so that the L -type's (IC) constraint is just binding: $\theta_L v(q_H) - p_H = \theta_L v(q_L^{LF}) - cq_L$. This implies that as the social planner increases q_H , the subsidy it offers becomes less generous because p_H must also increase, so the social planner assesses the marginal benefit of increasing q_H as $\pi_H \omega_H \theta_H v'(q_H) - \alpha \pi_H (c - \theta_L v'(q_H))$.

As a result, the optimal consumption level for the H -type is

$$q_H^* = (v')^{-1} \left(\frac{\alpha c}{\omega_H \theta_H - \alpha \theta_L} \right),$$

with corresponding price $p_H^* = \theta_L v(q_H^*) - \theta_L v(q_L^{LF}) + cq_L$, as illustrated in Figure 16.




Figure 16: Illustrating the optimal subsidized payment schedule with two types satisfying $\omega_L < \alpha < \omega_H$, with $\mathbf{E}[\omega] \leq \alpha$.

In short, the optimal subsidy mechanism with increasing welfare weights subsidizes consumption of $q_H^* > q_H^{LF}$ units and does not offer subsidies for lower consumption levels.

High Welfare Weights for Both Types: $\mathbf{E}[\omega] > \alpha$ and $\omega_H > \alpha$ We now turn to the final case with an active subsidy market, which is when $\mathbf{E}[\omega] > \alpha$ and $\omega_H > \alpha$.

In this case, the social planner finds it beneficial to subsidize the L -type and, if feasible, to further subsidize the H -type. The social planner can continue to offer subsidies to the H -type up until the point at which the L -type's (IC) constraint binds, so $p_H = \theta_L[v(q_H) - v(q_L)]$. As in the cases discussed in Appendix D.1, the social planner finds any non-distortive units of subsidy spending strictly beneficial, so the good is provided to the L -type for free. Substituting in the binding (IC) constraint, the social planner chooses q_L and q_H to maximize the weighted objective

$$\pi_L \omega_L \theta_L v(q_L) + \pi_H \omega_H [\theta_H v(q_H) - \theta_L [v(q_H) - v(q_L)]] + \alpha [\pi_H \theta_L [v(q_H) - v(q_L)] - c \pi_H q_H - c \pi_L q_L].$$

As a result, the optimal consumption levels are

$$q_L^* = (v')^{-1} \left(\frac{c \pi_L}{\pi_L \omega_L \theta_L + \pi_H \omega_H \theta_L - \alpha \pi_H \theta_L} \right), p_L^* = 0$$

while

$$q_H^* = (v')^{-1} \left(\frac{c}{\omega_H \theta_H - \omega_H \theta_L + \alpha \theta_L} \right), p_H^* = \theta_L [v(q_H^*) - v(q_L^*)].$$

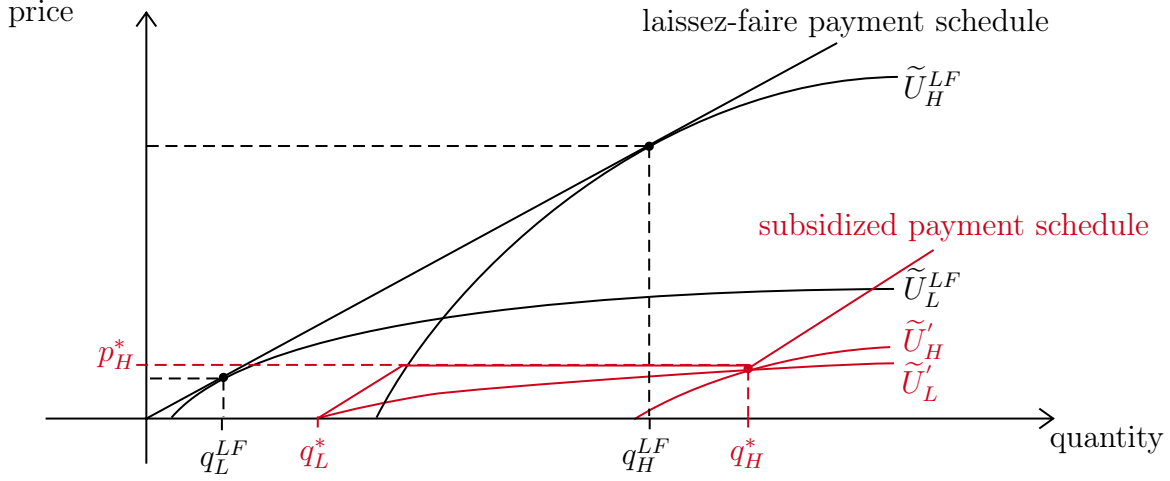


Figure 17: Illustrating an optimal subsidized payment schedule for $\mathbf{E}[\omega] > \alpha$ and $\omega_H > \alpha$ with upward distortion of both types' consumption.

The optimal mechanism is illustrated in Figure 17.

Key Takeaways

Summarizing our analysis, we have the following answers to the three questions posed at the start of this section: (i) the most cost-effective way to increase consumer surplus entails the least consumption distortion possible while satisfying the (NLS) and (IC) constraints; (ii) subsidies are optimal if and only if $\mathbf{E}[\omega] > \alpha$ or $\omega_H > \alpha$; and (iii) the optimal subsidy, in that case, distorts the consumption of one or both types upwards. Compared to the standard screening analysis on which this discussion was based, the subsidy design analysis is complicated by the possibility that either upward or downward incentive constraints may bind and the need to account for additional topping up deviations.

This paper generalizes this analysis to allow for richer consumer heterogeneity, as is needed to apply our results to real-world markets. While the key intuitions developed in this section extend, enriching our analysis to allow for more heterogeneity creates analytical complications not present in the two-type model. First, as we have seen, both upward and downward incentive constraints may bind, and while we could identify which binds for each type in the two-type model, doing so for the continuum model is more complex. Second, in the two-type model, nondistortive cash transfers are always available for sufficiently small cash subsidies; in the general model with a continuum of types, any subsidy offered to an interior type leads to some distorted consumption for nearby types. Third, in the two-type case, we have that welfare weights are either increasing

or decreasing, but in the general model, welfare weights may be nonmonotonic, providing an additional rationale for pooling adjacent types.

D.2 U- and Inverted U-Shaped Welfare Weights

In this section, we study the optimal subsidy mechanism under the assumption that the social planner’s welfare weights and the consumption type are not monotonic transformations of one another. Consequently, the good is neither “normal” nor “inferior” in the sense we described above. In particular, we focus on the case in which the social planner has U- and inverted U-shaped preferences, by which we mean the nonmonotone function $\omega(\theta)$ is quasiconcave or quasiconvex, respectively.

Inverted U-Shaped Preferences To characterize the optimal subsidy program under the assumption of inverted U-shaped welfare weights, we first determine the effect of that assumption on the distortion terms for the virtual welfare function, which depends on $\int_{\theta}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)$. As illustrated in Figure 18, the sign of $\omega(s) - \alpha$ changes sign at most twice, so the sign of the numerator of the distortion term changes at most twice as well.

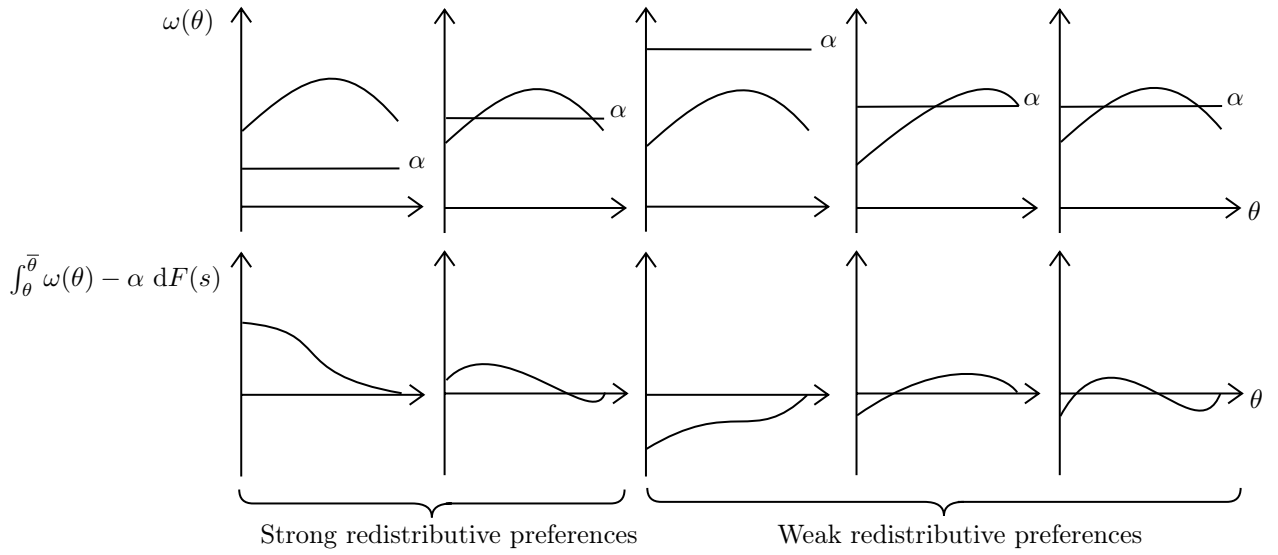



Figure 18: Plots of $\omega(\theta)$ and $\int_{\theta}^{\bar{\theta}} [\omega(s) - \alpha] dF(s)$ for inverted U-shaped welfare weights


When the cost of public funds is low (so $\mathbf{E}[\omega] > \alpha$), it is either (i) always positive, or (ii) positive and then negative. These distortion patterns are the same as those that apply to normal goods, discussed in Appendix B.1. As a result, the structure of the optimal subsidy mechanism

is the same as that case: a free endowment plus possibly additional quantity-dependent subsidies up to a capped level.

When the cost of public funds is high (with $\mathbf{E}[\omega] < \alpha$), there are three possibilities: (i) it is always negative, (ii) it is negative and then positive, or (iii) it is negative and then positive and then negative. In case (i), the distortion term is always negative (the same distortion pattern as in the normal good case with a high cost of public funds), and the social planner finds it optimal to implement the laissez-faire allocation rule with no subsidies. In case (ii), the social planner finds it optimal to distort higher levels of consumption upwards, as in the inferior goods case with a high cost of public funds, discussed in Appendix B.2. In case (iii), illustrated in the final panel of Figure 18, the social planner offers subsidies for *intermediate* consumption levels. That is, whereas the first units of the good are consumed at the competitive price, additional consumption is subsidized up to a capped level. That optimal mechanism, illustrated in Figure 19, combines the main features of the optimal subsidy mechanism for inferior goods (with a high cost of public funds) and the optimal subsidy mechanism for normal goods (with a low cost of public funds).



(a) Construction of the optimal subsidy mechanism (virtual welfare in red, lower bound in blue, subsidy type in black)



(b) Payment schedule (laissez-faire in blue, subsidized in black)

Figure 19: Optimal subsidy mechanism for inverted U-shaped preferences featuring subsidization of intermediate consumption levels

Two justification for studying inverted U-shaped welfare preferences are as follows. *First*, inverted U-shaped social preferences may arise naturally as a consequence of Downsian electoral competition (Downs, 1957). *Second*, inverted U-shaped social preferences may correspond to empirical regularities observed in the consumption data for the good, as follows. Suppose there are two goods, Good A and Good B, and that Good B is considered a close (superior) substitute to Good A. If total spending on the category of goods (Good A and Good B) is an increasing

but concave function of income, and the share of spending on Good A is a decreasing but concave function of income, then the resulting *total* spending on Good A is also a concave function of income, corresponding to inverted U-shaped preferences.

U-Shaped Preferences To characterize the optimal subsidy program under the assumption of U-shaped welfare weights, we again analyze the distortion terms for the virtual welfare function, which depends on $\int_{\bar{\theta}} [\omega(s) - \alpha] dF(s)$. Once again, as illustrated in Figure 20, the sign of $\omega(s) - \alpha$ changes sign at most twice, so the sign of the numerator of the distortion term changes at most twice as well.





Figure 20: Plots of $\omega(\theta)$ and $\int_{\bar{\theta}} [\omega(s) - \alpha] dF(s)$ for U-shaped welfare weights


When the cost of public funds is low, so $\mathbf{E}[\omega] > \alpha$, there are three possibilities: (i) the distortion term is positive, then negative, and then positive, (ii) the distortion term is positive and then negative, or (iii) the distortion term is always positive. Cases (ii) and (iii) involve the same distortion patterns as those for normal goods discussed in Appendix B.1, with a free allocation and possibly discounts for additional consumption. Case (i) is unlike other cases we have studied so far and results in allocations that are distorted upwards for both low and high types but undistorted for intermediate types, implemented via a free endowment (and possibly additional discounts for low levels of consumption) and additional discounts for consumption beyond a certain level. We illustrate the optimal subsidy mechanism in Figure 21.

When the cost of public funds is high, so $\mathbf{E}[\omega] < \alpha$, there are two possibilities: (i) the distortion term is always negative, or (ii) the distortion term is negative and then positive. These distortion

patterns resemble those associated with inferior goods with a high cost of public funds, so the optimal subsidy mechanisms in these cases are the same as those discussed in Appendix B.2.



(a) Construction of the optimal subsidy mechanism (virtual welfare in red, lower bound in blue, subsidy type in black)



(b) Payment schedule (laissez-faire in blue, subsidized in black)

Figure 21: Optimal subsidy mechanism for U-shaped preferences featuring a free endowment and subsidization of higher consumption levels

D.3 Unit Demand

In this section, we discuss the optimal subsidy mechanism under the assumption of unit demand. One minor complication that arises in this case is that the valuation function is not strictly concave, and as a result, the optimal subsidy allocation rule is not generally unique. In this section, we characterize the set of all optimal subsidy mechanisms in the unit demand case.

For the purpose of this section, we let $v : [0, 1] \rightarrow [0, 1]$ be defined by $v(q) = q$, where q is interpreted as the *probability* that the consumer is allocated the good.²⁴ We suppose, moreover, that $c \in \text{int } \Theta$, so not all agents purchase the good in the laissez-faire economy.

The optimal subsidy mechanism satisfies the *linear* program

$$\max_{q: \Theta \rightarrow [0,1]} \int_{\Theta} \alpha [J(\theta) - c] q(\theta) dF(\theta) + (\text{terms independent of } q)$$

such that q is non-decreasing and satisfies (LB): $q(\theta) = 1$ for all $\theta \geq c$.

²⁴ Although we show in the divisible good case that the restriction to deterministic mechanisms is without loss of generality, that argument (in Appendix C.1) relies on the divisibility of the good (so that the certainty equivalent is a feasible allocation). In the unit demand setting, we restore divisibility by treating q as the probability of assignment of the good, but as a result—the mechanisms we study are necessarily randomized mechanisms.

A characterization of the set of optimal subsidy mechanisms is contained in Proposition 16

Proposition 16 (optimal subsidies with unit demand). *There exists an optimal subsidy mechanism in the unit demand case which is deterministic (that is, with $q(\theta) \in \{0, 1\}$ for all $\theta \in \Theta$), and any such deterministic subsidy mechanism allocates the good to all types higher than the cutoff $\hat{\theta}$ satisfying*

$$\hat{\theta} \in \arg \max_{\theta \in [\underline{\theta}, c]} \int_{\theta}^c [J(s) - c] dF(s).$$

In particular, if $\hat{\theta} = \underline{\theta}$, subsidies require $\mathbf{E}[\omega] > \alpha$, and the subsidized price of the good is zero. Otherwise, the subsidized price of the good is $p = \hat{\theta}$ and p satisfies

$$c = p + \frac{\int_p^{\bar{\theta}} \omega(s) - \alpha dF(s)}{\alpha f(p)}.$$

Under the following two conditions, there also exist optimal subsidy mechanisms with randomization:

- (i) $\int_c^{\theta} [J(s) - c] dF(s) \leq 0$ for all $\theta \geq c$, and
- (ii) $\int_{\theta}^c [J(s) - c] dF(s) \leq 0$ for all $\theta \leq c$, with equality for some $\theta < c$.

In that case, writing $\theta_1 < \theta_2 < \dots < \theta_I := c$ for the set of types less than c for which $\int_{\theta}^c [J(s) - c] dF(s) = 0$, any optimal randomized subsidy mechanism satisfies $q^(\theta) = 0$ for any $\theta < \theta_1$, $q^*(\theta) = 1$ for any $\theta > c$, and $q^*(\theta) = q_i$ on any interval (θ_i, θ_{i+1}) for $i = 1, 2, \dots, I - 1$, where $\{q_i\} \subseteq [0, 1]$ is any nondecreasing sequence of constants. In that case, the price p_i for purchasing the good with probability q_i can be calculated as $p_1 = q_1 \theta_1$, and $p_i = p_{i-1} + (q_i - q_{i-1}) \theta_i$.*

Proof. The characterization of the optimal allocation rule described in Proposition 16 follows as a direct application of the solution of general linear programs with lower-bound constraints, presented in Proposition 17 below. The two outstanding details concern the price of the subsidized good and are as follows:

- (a) The claim that $c = p + \frac{\int_p^{\bar{\theta}} \omega(s) - \alpha dF(s)}{\alpha f(p)}$ for any optimal subsidy price $p < c$. This is the first-order condition for the maximization problem $\max_{p \leq c} \int_p^c [J(\theta) - c] dF(\theta)$, which is a necessary condition for an interior maximum.
- (b) The claim that free provision of the good is never optimal when the opportunity cost of public funds is high. To see this, note that when the opportunity cost of public funds is

high, $J(\underline{\theta}) < \underline{\theta} < c$, and thus $\int_p^c [J(\theta) - c] dF(\theta)$ is strictly increasing at $p = \underline{\theta}$. As a result, $\underline{\theta}$ cannot be the maximizer.

□

General Linear Programs with Lower-Bound Constraints In this section, we solve the linear program

$$\max_{q: \Theta \rightarrow [0, A]} \int_{\Theta} \phi(\theta) q(\theta) dF(\theta), \text{ subject to } \forall \theta \in \Theta : q(\theta) \geq q^L(\theta) \quad (\text{LP})$$

for a monotone function $q^L : \Theta \rightarrow [0, A]$ and an integrable $\phi : \Theta \rightarrow \mathbb{R}$, assumed to be nonzero almost everywhere in Θ . While this result can be obtained by taking appropriate limits in Theorem 2' above, the statement and proof of this special case is revealing in its own right, and so we include it here.

In this section, moreover, we assume that $A = 1$ (this is without loss, as we can always divide the quantity by A and think of q as the fraction of the total A allocated). In that case q and q^L may be interpreted as the distribution functions of random variables (adjusting the value of q and q^L at $\bar{\theta}$, if necessary, which clearly has no impact on the value of the linear program), and the lower-bound constraint is a *first-order stochastic dominance constraint*, requiring that q^L dominate q in the first-order stochastic dominance sense.

For the statement of this result, write $\Phi(\theta) = \int_{\underline{\theta}}^{\theta} \phi(s) dF(s)$ and let $\text{decmin } \Phi$ be the *decreasing minorant* of Φ , which is the largest nonincreasing function everywhere below Φ , as illustrated in Figure 22. The type space Θ can be partitioned into intervals Θ^i (intersecting only at their endpoints and increasing in the strong set order) in which either (a) $\Phi(\theta) = \text{decmin } \Phi(\theta)$ for all $\theta \in \Theta^i$, or (b) $\Phi(\theta) > \text{decmin } \Phi(\theta)$ for all $\theta \in \text{int } \Theta^i$ (with equality at the endpoints). Write $\Theta = \cup_i \Theta^i$ for such a partition, and note that there may be adjacent intervals of type (b), as in Figure 22.

The solutions of (LP) are characterized in the following proposition.

Proposition 17 (solutions to linear program). *Any solution of (LP) satisfies the following:*

- (i) *If $\Phi(\theta) = \text{decmin } \Phi(\theta)$ on Θ^i , then $q^*(\theta) = q^L(\theta)$ for all $\theta \in \Theta^i$.*
- (ii) *If $\Phi(\theta) > \text{decmin } \Phi(\theta)$ on $\text{int } \Theta^i$ and $\sup \Theta^i = \bar{\theta}$, then $q^*(\theta) = 1$ for all $\theta \in \Theta^i$.*




Figure 22: Constructing the Decreasing Minorant

(iii) Otherwise, if $\Phi(\theta) > \text{decmin } \Phi(\theta)$ on $\text{int } \Theta^i$, write $\theta_+ = \sup \Theta^i$. Then $q^*(\theta)$ is constant on $\text{int } \Theta^i$, with that constant determined as (a) any number in $[\lim_{\theta \uparrow \theta_+} q_L(\theta), q_L(\theta_+)]$ if Θ^{i+1} satisfies $\text{decmin } \Phi(\theta) = \Phi(\theta)$ on Θ^{i+1} , or (b) any number in $[\lim_{\theta \uparrow \theta_+} q_L(\theta), q^*(\theta_+)]$ otherwise.

While the statement of Proposition 17 appears complicated, the intuitive description of the optimizers is simpler: wherever $\Phi(\theta) = \text{decmin } \Phi(\theta)$, any optimizer equals the lower bound, otherwise the optimizer is a constant, typically determined by the value of q^L at the right endpoint of that interval. The complication arises from the possibility that q^L is discontinuous at that right endpoint (in which case the optimizer can be constant at any level within that jump) and the possibility that there are adjacent intervals on which $\Phi > \text{decmin } \Phi$, in which case the feasible constants on each interval depend on the constant chosen on the subsequent interval.

Proof. We have the following upper bound on the value of (LP):

$$\begin{aligned}
\int_{\Theta} \phi(\theta) q(\theta) dF(\theta) &= q(\bar{\theta})\Phi(\bar{\theta}) - \int_{\Theta} \Phi(\theta) dq(\theta) \\
&\leq \Phi(\bar{\theta}) - \int_{\Theta} \text{decmin } \Phi(\theta) dq(\theta) \\
&\leq \Phi(\bar{\theta}) - \int_{\Theta} \text{decmin } \Phi(\theta) dq^L(\theta) \\
&= \Phi(\bar{\theta}) - \text{decmin } \Phi(\bar{\theta}) + \int_{\Theta} \frac{d}{d\theta} \text{decmin } \Phi(\theta) q^L(\theta) d\theta.
\end{aligned}$$

The first and fourth lines use integration-by-parts for the Lebesgue-Stieltjes Integral, the second line follows by the construction of $\text{decmin } \Phi$ as a minorant (and uses the fact that $q(\bar{\theta}) = 1$ because $q^L(\bar{\theta}) = 1$ by assumption), and the third by the fact that q^L first-order stochastically dominates q and $-\text{decmin } \Phi$ is increasing.

We now argue that the upper bound is obtained by any solution taking the form described in Proposition 17 and that any other feasible q results in a strictly lower objective value. We calculate $\int_{\Theta} \phi(\theta)q^*(\theta) dF(\theta)$ separately for each interval of type (i), (ii), and (iii) described in Proposition 17.

On an interval Θ^i of type (i), we have

$$\int_{\Theta^i} \phi(\theta)q^*(\theta) dF(\theta) = \int_{\Theta^i} \phi(\theta)q^L(\theta) dF(\theta) = \int_{\Theta^i} \frac{d}{d\theta} \text{decmin } \Phi(\theta)q^L(\theta) d\theta,$$

because $q^*(\theta) = q^L(\theta)$ on Θ^i and because $\text{decmin } \Theta(\theta) = \Theta(\theta)$. On the other hand, for any feasible q which strictly exceeds q^L on some positive measure subset of Θ^i , it is clear that the objective value is strictly lower because $\phi(\theta) < 0$ by construction on those intervals.

On any interval Θ^i of type (ii), we have

$$\int_{\Theta^i} \phi(\theta)q^*(\theta) dF(\theta) = \int_{\Theta^i} \phi(\theta) dF(\theta) = \Phi(\bar{\theta}) - \text{decmin } \Phi(\bar{\theta}) + \int_{\Theta^i} \frac{d}{d\theta} \text{decmin } \Phi(\theta)q^L(\theta) d\theta,$$

by construction of $\text{decmin } \Phi$ (noting that the integral term is zero, as $\text{decmin } \Phi$ is constant). On the other hand, for any other feasible q , we have

$$\int_{\Theta^i} \phi(\theta)q(\theta) dF(\theta) = \Phi(\bar{\theta}) - \Phi(\theta_-) - \int_{\Theta^i} (\Phi(\theta) - \Phi(\theta_-)) dq(\theta),$$

where θ_- is the left endpoint of Θ^i so $\text{decmin } \Phi(\bar{\theta}) = \Phi(\theta_-)$ by construction, which is clearly less than the result obtained above (because $\Phi(\theta) - \Phi(\theta_-) > 0$ and q is a positive measure).

Finally, on any interval of type (iii), we have

$$\int_{\Theta^i} \phi(\theta)q^*(\theta) dF(\theta) = \int_{\Theta^i} \phi(\theta) dF(\theta) = 0 = \int_{\Theta^i} \frac{d}{d\theta} \text{decmin } \Phi(\theta)q^L(\theta) d\theta, \Phi$$

by construction of $\text{decmin } \Phi$ and because $\frac{d}{d\theta} \text{decmin } \Phi$ is constant on any such interval. Once again, for any other feasible q , we have

$$\int_{\Theta^i} \phi(\theta)q(\theta) dF(\theta) = - \int_{\Theta^i} (\Phi(\theta) - \Phi(\theta_-)) dq(\theta) < 0.$$

Adding these expressions together, we obtain that $\int_{\Theta} \phi(\theta)q^*(\theta) dF(\theta)$ obtains the upper bound we derived above, and that the objective value for any alternative q is strictly worse. □