Optimal Redistribution Through Subsidies*

Zi Yang Kang[†] Mitchell Watt[‡]

This version: October 25, 2024
Latest version: available **here**

Abstract

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 Dworczak, 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 wishes to redistribute in a market by offering subsidies to consumers. Subsidies can be nonlinear: the social planner can adjust the marginal subsidy for each additional unit of the good consumed. Subsidies are also costly. When should the social planner offer any subsidy at all? If she chooses to offer a subsidy, how should it be designed?

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 a wide range of goods, including child care, education, energy, fuel, health care, and transportation. These subsidies are often nonlinear by design. For instance, food stamp programs—such as the Supplemental Nutrition Assistance Program (SNAP) in the United States—provide a full subsidy for a fixed monetary amount of food, but zero subsidy thereafter.

In this paper, we study the problem of optimal redistribution through subsidies. In our model, the social planner can subsidize consumption by reimbursing some or all of the private market price for different quantities of the good, while consumers can always purchase additional units at the private market price. 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 presence of the private market constrains the social planner's ability to redistribute. The social planner would like to screen consumers by setting higher marginal prices for consumption levels chosen by low-need consumers. However, because consumers can always access the private market, the marginal cost of any unit of consumption cannot exceed the private market price (or equivalently, the total subsidy received by consumers must be increasing in the quantity consumed), limiting the extent of redistribution from low-need to high-need consumers. Using mechanism design—and particularly the taxation principle (Guesnerie, 1981)—we reformulate this constraint, showing it to be equivalent to requiring that each consumer's total consumption exceeds his laissez-faire consumption level.

In our first main result, Theorem 1, we identify a sufficient statistic that determines whether the social planner strictly benefits from offering subsidies. Specifically, the social planner optimally intervenes using 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 offers to reimburse a small amount of spending for consumers who purchase more than the laissez-faire demand of that type, the increase in their weighted 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.

To better understand the constraints imposed by 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 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 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.

Our second main result, Theorem 2, provides an explicit characterization of 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 the condition in Theorem 1 is satisfied, the *marginal* unit purchased by that type in the optimal subsidy mechanism is at least partially subsidized.

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

[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.

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 consumers from topping up 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 consumers from topping up 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 larger 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 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 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

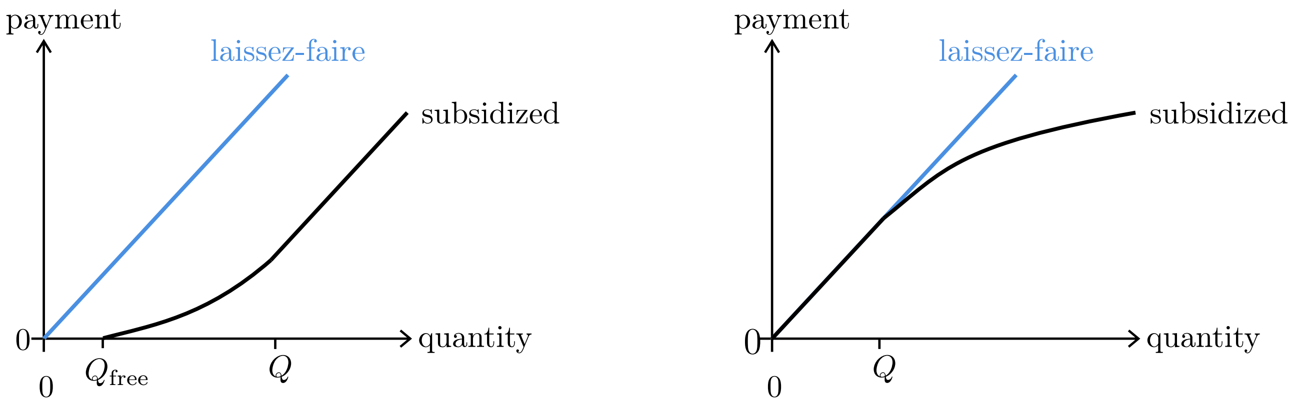

(a) Optimal payment schedule when welfare weights are negatively correlated with demand and, on average, higher than the opportunity cost of funds.

(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.

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 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 over the history of the program 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 topping up feature, 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 that accounts 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 contributes to a literature on in-kind transfers in public finance (surveyed by Currie and Gahvari (2008)) by deriving the optimal form of in-kind transfers (namely, vouchers or direct provision) as a function of 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 topping up 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, Dworczak 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 ⓡ Kominers ⓡ Akbarpour, 2021; Akbarpour ⓡ Dworczak ⓡ 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 ⓡ Budish ⓡ Dworczak ⓡ 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 rather than the possibility of topping up, 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 have access different production technologies for the sake of tractability, our paper develops tools to analyze when the social planner and the private market to have access to the same production technology (which we interpret as the ability of the social planner to contract public production out to private firms or to directly subsidize private purchases).

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 in-kind subsidy design for a redistributive social planner in a setting with topping up. In Section 3, we explore the constraints imposed by topping up and discuss 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 formulate a model of subsidy design for a social planner with redistributive preferences and establish a laissez-faire benchmark for comparison with the subsidy designs.

2.1 Economic Primitives

There is a unit mass of risk-neutral consumers in a market for a divisible good.

Each consumer is privately informed of his type θ, distributed according to an absolutely continuous cumulative distribution function F with positive density function f supported on the type space $\Theta := [\underline{\theta}, \overline{\theta}]$, where $\underline{\theta} > 0$. Consumer preferences are quasilinear in money, with each consumer's type determining his value from consumption: a consumer of type θ derives utility $\theta v(q) - t$ from consuming q units of the good and paying t units of money. Unless otherwise stated,[2] we assume that the valuation function $v : [0, A] \to \mathbb{R}_+$ is twice continuously differentiable, with $v' > 0$ and $v'' < 0$. For notational simplicity, we extend the domain of $(v')^{-1}$ to the entire real line, setting $(v')^{-1}(z) = 0$ for $z \geq v'(0)$ and $(v')^{-1}(z) = A$ for $z \leq v'(A)$, allowing us to write the individual demand curve as

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

For now, we assume that the good is supplied competitively by producers in a private market, each with a constant marginal cost of production, denoted by $c > 0$. In Section 6.1, we discuss extensions of our results that allow alternative private market structures.

[2] In Appendix D.3, we discuss the important special case of unit demand, for which the valuation function $v(q) = q$ is not strictly concave.

2.2 Laissez-Faire Equilibrium

In the laissez-faire equilibrium of the market, the good is priced at its marginal cost, c. Given that price, each consumer solves his utility maximization problem,

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

The strict concavity and monotonicity of v implies that the solution to this optimization problem is unique for each type and increasing in θ. We call that solution the *laissez-faire allocation rule* and denote it by $q^{\mathrm{LF}}(\theta) = D(c, \theta)$. The First Welfare Theorem (Arrow, 1951) implies that q^{LF} maximizes utilitarian surplus, which is the sum of consumers' utilities minus production costs.

2.3 Social Planner's Objective

Instead of maximizing utilitarian surplus, the social 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 consumer surplus, which is the increase in expected social welfare associated with giving a dollar to a consumer of type θ.[4] We assume that $\omega : \Theta \to \mathbb{R}_+$ 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:

(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.

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

(b) *Cost of Spending*: The social planner assigns a weight α to the cost of spending, capturing the *opportunity cost of funds*. 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]

Putting these components together, the social planner's objective is *weighted total surplus*,

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

where $q(\theta)$ is the quantity consumed by a consumer of type θ and $t(\theta)$ is the consumer's payment to the social planner. The utilitarian objective is obtained by setting $\omega(\theta) \equiv \alpha$. When the welfare weights are not all identical, the social planner would like to redistribute surplus towards consumers with $\omega > \alpha$ and away from consumers with $\omega < \alpha$.

2.4 Subsidy Design

To reallocate goods and money among different types of consumers, the social planner can subsidize purchases of the good.[6] To do so, the social planner announces a *subsidized payment schedule* $T^{\mathrm{S}} : [0, A] \to \mathbb{R}_+$, where $T^{\mathrm{S}}(z)$ is the total payment that a consumer must make to obtain z units of the good from the social planner, so $S(z) = cz - T^s(z)$ is the *total subsidy* received by a consumer who purchases z units of the good. While the private market's payment schedule is linear under the assumption of perfect competition, the social planner's subsidy schedule may involve nonlinear pricing. We allow the social planner to give away some goods for free by setting

[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.

[6] For now, we assume that the full unit mass of consumers is eligible for the subsidy program, and we discuss the implications of our results for a model in which the social planner can determine eligibility rules based on observable characteristics in Section 7.

$T^{\mathrm{S}}(z) = 0$ for some $z > 0$, but require $T^{\mathrm{S}}(z) \geq 0$, so she cannot reimburse more than the laissez-faire cost of the good to consumers. We require that $T^{\mathrm{S}}(0) = 0$ so the consumer can always opt out of the subsidized market. The social planner shares the same production technology as the private market, with the same marginal cost c for the good.

A key assumption, introduced in this paper, is that consumers can *top up* their consumption of subsidized units of the good by purchasing additional units at the competitive market price of c.[7] We assume that the social planner can verify that the consumer has consumed at least the number of units of the good he reports to have consumed in order to receive the subsidy but cannot prevent the consumer from purchasing additional units of the good beyond the reported amount. As a result, consumers cannot access the subsidized payment schedule multiple times nor can they resell quantities of the good.[8]

Mechanism design formulation Instead of focusing directly on the subsidized payment schedule, we frame the social planner's choice as a mechanism design problem. The social planner determines each consumer's *total* consumption, including both subsidized and privately purchased units. Since the social planner has the same production technology as the private market, total consumption is sufficient to determine total costs and surplus. In the mechanism design formulation, the social planner chooses a direct mechanism (q, t), consisting of:

(a) an allocation rule $q : \Theta \to [0, A]$, where $q(\theta)$ is the total quantity consumed by type θ, and

(b) a payment rule $t : \Theta \to \mathbb{R}$, where $t(\theta)$ is the total payment made by a consumer of type θ.

While this formulation implicitly restricts attention to deterministic mechanisms, we show in Appendix C.1 that this is without loss: assuming that v is strictly concave and the good is divisible, the social planner never benefits from randomization.

We now specify which mechanisms are *feasible*. Without loss of generality, we focus on direct mechanisms, in which consumers truthfully report their types to the social planner, who then determines their allocations and payments. The possibility of topping up in the private market

[7] In Kang and Watt (2024), we study the alternative case in which the consumer is precluded from topping up his consumption of the good and must choose between a public option or a private market.

[8] Resale may be a concern in some subsidy markets. Our model could also be extended to include costly resale, leading to a lower bound on the marginal price in the total payment schedule, but we focus on the case without resale to isolate the effect of the private market on subsidy design. Empirical evidence from the US SNAP program finds resale (a form of food stamp fraud) to be a second-order concern in practice, with Woody, Zhao, Lund and Wu (2024) estimating that less than 1.5% of food stamp spending is fraudulent.

introduces new potential deviations by consumers, leading to stronger constraints on the mechanism design problem.

The first is a strengthening of the standard incentive compatibility constraint, ruling out not only type misreports but also topping up in the private market. Let $V(\theta, \hat{\theta})$ be the maximum utility a consumer of type θ could obtain by reporting type $\hat{\theta}$ and topping up his consumption, so

$$V(\theta, \hat{\theta}) = \max_{z \geq q(\hat{\theta})} \left\{ \theta v(z) - t(\hat{\theta}) - c(z - q(\hat{\theta})) \right\},$$

The *incentive compatibility constraint with topping up* requires that

$$\text{for all } \theta \in \Theta, \qquad \theta v(q(\theta)) - t(\theta) = \max_{\theta' \in \Theta} V(\theta, \theta'). \tag{IC-T}$$

Second, a consumer might bypass the subsidized market entirely and rely solely on the private market, leading to the *individual rationality constraint*:

$$\text{for all } \theta \in \Theta, \qquad \theta v(q(\theta)) - t(\theta) \geq U^{\text{LF}}(\theta). \tag{IR}$$

Finally, we assume the social planner is restricted from setting a negative payment for any quantity of the goods, so it cannot make lump-sum transfers to consumers via the subsidy market. This leads to the *no lump-sum transfers constraint*:

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

This rules out lump-sum transfers within the mechanism but not outside the mechanism, which may influence the social planner's preferences for spending within the mechanism by changing the opportunity cost of funds, α.

In Figure 2, we illustrate the constraints in the mechanism design problem in price-quantity space.[9] For a consumer of type θ, the (NLS) constraint requires that the price-quantity pair $(t(\theta), q(\theta))$ lies above the quantity axis, while the (IR) constraint requires that it lies below θ's laissez-faire indifference curve (which is tangent to the price schedule at $q^{\text{LF}}(\theta)$). Fixing such a price-quantity pair—say $(t(\theta), q(\theta))$, the red point in Figure 2—the (IC-T) constraint imposes several conditions on the price-quantity pair assigned to a type $\hat{\theta} > \theta$, corresponding to different possible deviations by consumers. First, as in standard screening problems (e.g., Stiglitz, 1982;

[9] In Appendix D.1, we extend the analysis of Figure 2 to calculate the optimal subsidy mechanism in an example with two types of consumers.

Maskin and Riley, 1984), to prevent either type preferring the other's allocation, $(t(\hat{\theta}), q(\hat{\theta}))$ must lie above θ's indifference curve through $(t(\theta), q(\theta))$ and below $\hat{\theta}'s$ indifference curve through the same point. Second, unlike standard screening problems, the (IC-T) constraint requires that the marginal price difference $t(\hat{\theta}) - t(\theta)$ be lower than the market price for additional consumption $c[q(\hat{\theta}) - q(\theta)]$, so $(t(\hat{\theta}), q(\hat{\theta}))$ must lie below the blue line in Figure 2. Together, this means that $(t(\hat{\theta}), q(\hat{\theta}))$ must lie in the shaded region of Figure 2.

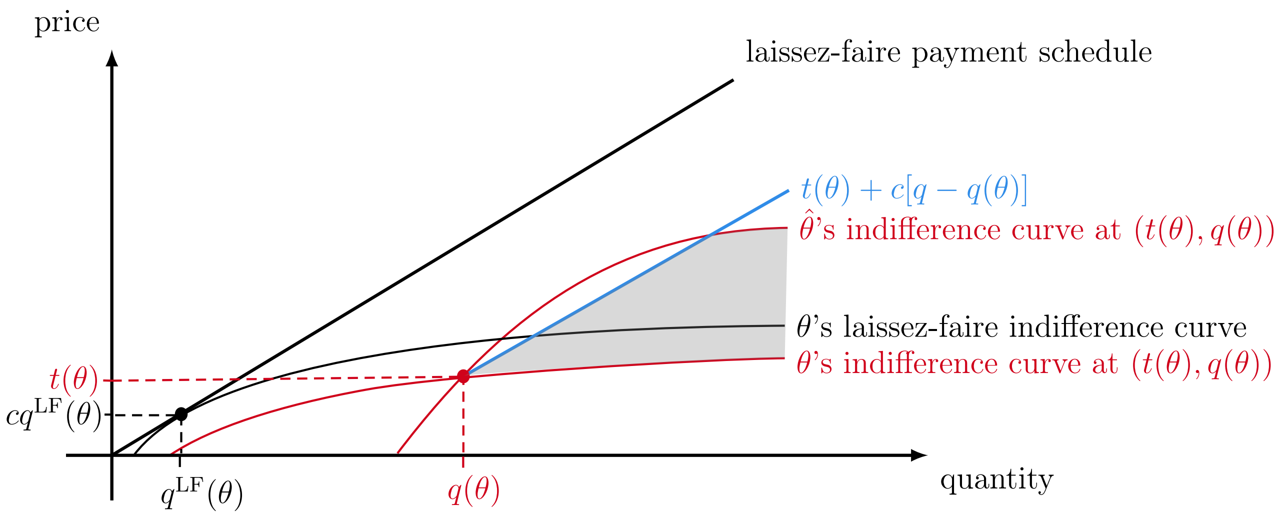

Figure 2: Illustrating the constraints in price-quantity space.

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

$$\max_{(q,t)} \int_{\Theta} \Big\{\omega(\theta) \underbrace{[\theta v(q(\theta)) - t(\theta)]}_{\text{consumer surplus}} - \alpha \underbrace{[cq(\theta) - t(\theta)]}_{\text{total costs}} \Big\} \, \mathrm{d}F(\theta), \tag{SUB}$$

such that (q, t) satisfies (IR), (IC-T) and (NLS).

Because the laissez-faire mechanism $(q^{\mathrm{LF}}, cq^{\mathrm{LF}})$ is always feasible, the set of feasible mechanisms is always nonempty, and the optimization program (SUB) 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.

Implementation We say that a mechanism (q, t) is *implemented* by a subsidized payment schedule T^{S} if

$$\text{for all } \theta \in \Theta, \quad q(\theta) \in \underset{z\in[0,A]}{\arg\max} \left[\theta v(z) - T^{\mathrm{S}}(z)\right] \text{ and } t(\theta) = T^{\mathrm{S}}(q(\theta)).$$

Because the social planner is equally efficient as the private market, there is no difference in the model between *(a)* the social planner producing all units of the good and offering them for sale at the subsidized payment schedule T^{S} and *(b)* all goods being produced and sold in the private market, with the social planner paying total subsidies $S(z)$ for z units of consumption (or some mix of the two).

Note that any feasible mechanism can be implemented by a subsidized payment schedule $T : [0, A] \to \mathbb{R}$ as follows. For any $z \in \operatorname{im} q$, let $T(z) = t(\theta)$ for any $\theta \in q^{-1}(z)$. Then, for any $z \notin \operatorname{im} q$, set $T(z)$ equal to the payment for the least $z' > z$ with $z' \in \operatorname{im} q$, or $+\infty$ if there is no such z'.

2.5 Model Discussion

The key assumption introduced in this paper is that consumers can top up subsidized allocations in a private market. In many real-world subsidy markets, including those for food, transport and fuel, the prohibitive costs associated with monitoring and enforcing private market restrictions makes topping up practically unavoidable. By assuming that consumers can buy at the competitive price c, the model rules out the social planner setting a tax on private market purchases. Without this assumption, the social planner could set an infinite tax on the private market and implement any (nondecreasing) allocation via 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 6.2.

Another important assumption in our subsidy design model is the no lump-sum transfers constraint, (NLS). In practice, this constraint 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 the production technology facing the social planner. 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.[10] 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 (for both the public and private markets) is perfectly inelastic. 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 better understand the effects of the topping up constraint in isolation, 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 of 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.

Theorem 1 (scope of optimal subsidy). *The optimal mechanism (q^*, t^*) strictly improves on the laissez-faire outcome if and only if there exists a type $\hat{\theta} \in \Theta$ for which*

$$\mathbf{E}_\theta\left[\omega(\theta) \,|\, \theta \geq \hat{\theta}\right] > \alpha.$$

[10] 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).

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 $\max \omega > \alpha$.

To obtain this result, we first study the implications of the (IC-T) constraint for the subsidized payment schedule. We show that the consumer's ability to top up requires the social planner to set a payment schedule with total subsidies increasing with the quantity consumed. This feature aligns with the goal of redistributing to higher welfare weight consumers when ω is increasing in θ, but conflicts with this objective when ω is decreasing. As a result, there is a narrower scope for redistribution using in-kind subsidies for goods with negative correlation between demand and need than for positive correlation. More generally, if the average welfare weight of consumers with type at least $\hat{\theta}$ exceeds the opportunity cost of funds, we show that the social planner can use nonlinear subsidies to benefit those consumers more than the associated costs of distorted consumption for consumers with lower types.

3.1 Proof of Theorem 1

Reformulating the (IC-T) Constraint. The (IC-T) constraint, as formulated above, imposes joint restrictions on allocations and prices in any feasible subsidy program. To better understand the constraints on redistribution imposed by the consumers' ability to top up in a private market, we now study the implications of (IC-T) for subsidized price schedules and allocation rules.

Proposition 1 (reformulating the incentive constraint with topping up)**.** *Let (q, t) satisfy the (IR) and (NLS) constraints. The following are equivalent:*

(i) (q, t) satisfies the (IC-T) constraint.

(ii) (q, t) can be implemented by a subsidized payment schedule $T^{\mathrm{S}} : [0, A] \to \mathbb{R}_+$ such that the total subsidy $S(z) = cz - T^{\mathrm{S}}(z)$ is positive and nondecreasing.

(iii) (q, t) satisfies the standard incentive compatibility constraint

$$\text{for all } \theta \in \Theta, \qquad \theta \in \arg\max_{\theta'} \left[\theta v(q(\theta')) - t(\theta')\right], \tag{IC}$$

and the lower-bound constraint

$$\textit{for all } \theta \in \Theta, \qquad q(\theta) \geq q^{\mathrm{LF}}(\theta). \tag{LB}$$

We prove Proposition 1 in Appendix C.3, but now briefly discuss the intuition.

The fact that *(i)* implies *(ii)* follows by contradiction. If total subsidies $S(q)$ decreased between $q(\theta)$ and $q(\theta')$ with $q(\theta') > q(\theta)$, a consumer of type θ' would prefer to misreport his type as θ and buy the difference $q(\theta') - q(\theta)$ in the private market, violating the (IC-T) constraint. This can be seen in Figure 2 above: the (IC-T) constraint from allocation $(t(\theta), q(\theta))$ lowers the set of feasible prices for $q \geq q(\theta)$ by the vertical gap between the laissez-faire payment schedule and $(t(\theta), q(\theta))$, which is the total subsidy received by type θ. Therefore, the total subsidy for higher consumption levels must be at least as large. Conversely, if the payment schedule has slope at most c (as implied by *(ii)*), a consumer of type θ has no incentive to report $\theta' \neq \theta$ to receive $q(\theta')$ and then top up at the market price c, because the total payment would exceed the subsidized price for the same quantity.

To see that *(i)* implies *(iii)*, note that the (IC) constraint is implied directly by (IC-T), and the lower-bound constraint (LB) holds because any consumer allocated less than his laissez-faire demand would purchase additional units in the private market to reach that consumption level. The converse follows from the Milgrom and Segal (2002) envelope theorem implied by (IC), which requires that the marginal price at $q(\theta)$ units of consumption equals the marginal value of consumption to type θ, which is always less than c when (LB) is satisfied. This implies that the subsidized payment schedule has a slope at most c, so the total subsidy is increasing in the consumption level, completing a chain of implications from *(iii)* to *(ii)* to *(i)*.

Applying Proposition 1. The requirement that total subsidies increase with the quantity consumed means that the social planner's redistributive power depends critically on the relationship between welfare weights and consumer demand. In particular, while Proposition 1 is stated in terms of $T^{\mathrm{S}}(z)$, it also implies that total subsidies are increasing in θ since consumer demand is increasing in θ under any subsidized payment schedule by the supermodularity of the consumer's objective, $\theta v(z) - T^{\mathrm{S}}(z)$.

When the social planner assigns higher welfare weights to low-demand consumers, meaning $\omega(\theta)$ is decreasing, the monotonicity of total subsidies conflicts with the social planner's redistributive preferences. By Theorem 1, the social planner offers in-kind subsidies only when $\mathbf{E}[\omega] > \alpha$. On the other hand, whenever $\mathbf{E}[\omega] > \alpha$, the (NLS) binds because transferring a dollar to all consumers would increase weighted surplus by $\mathbf{E}[\omega]$ but cost α, which is profitable if and only if $\mathbf{E}[\omega] > \alpha$. As a result, when demand for the good and need are negatively correlated, the social planner only subsidizes consumption if she is restricted to redistributing in-kind and would otherwise prefer to offer cash transfers to all consumers.

On the other hand, to see why the social planner does not offer in-kind subsidies when ω is decreasing and $\mathbf{E}[\omega] \leq \alpha$, 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. The social planner's benefit of a subsidy schedule S, measured by the expected increase in weighted consumer surplus, is bounded as follows:

$$\begin{aligned}\int_\Theta \left[\omega(\theta)(U(\theta) - U^{\mathrm{LF}}(\theta))\right] \, \mathrm{d}F(\theta) &\leq \int_\Theta \left[\omega(\theta)S(q(\theta))\right] \, \mathrm{d}F(\theta) \\ &\leq \mathbf{E}_\theta[\omega(\theta)]\, \mathbf{E}_\theta[S(q(\theta))],\end{aligned}$$

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). 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.[11] Since $\mathbf{E}[\omega] \leq \alpha$, the social planner prefers not to make lump-sum transfers and, thus, does not opt for the more regressive subsidy.

When the social planner has higher welfare weights for high-demand consumers, so $\omega(\theta)$ is increasing, the monotonicity of total subsidies aligns with the social planner's redistributive preferences. In this case, Theorem 1 shows that the social planner redistributes whenever there is a consumer type with $\omega(\theta) > \alpha$. More generally, whenever $\mathbf{E}[\omega|\theta \geq \hat{\theta}] > \alpha$ for some $\hat{\theta} \in \Theta$ (which necessarily holds for $\hat{\theta}$ near $\overline{\theta}$ when ω is decreasing and $\max \omega > \alpha$), we can construct a subsidized payment schedule that increases the social planner's payoff compared to laissez-faire.

To establish Theorem 1, consider the subsidized payment schedule illustrated in Figure 3, where the social planner reimburses ε of spending for any consumer buying more than $q^{\mathrm{LF}}(\hat{\theta})$ 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 $\int_{\hat{\theta}}^{\overline{\theta}} [\omega(\theta) - \alpha] \, \mathrm{d}F(\theta)$, which

[11] 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.

is linear in ε. On the other hand, the (IC-T) constraint requires the social planner to increase consumption for consumers just below $\hat{\theta}$ to $q^{\mathrm{LF}}(\hat{\theta})$, as shown in Figure 3. 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,[12] 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.

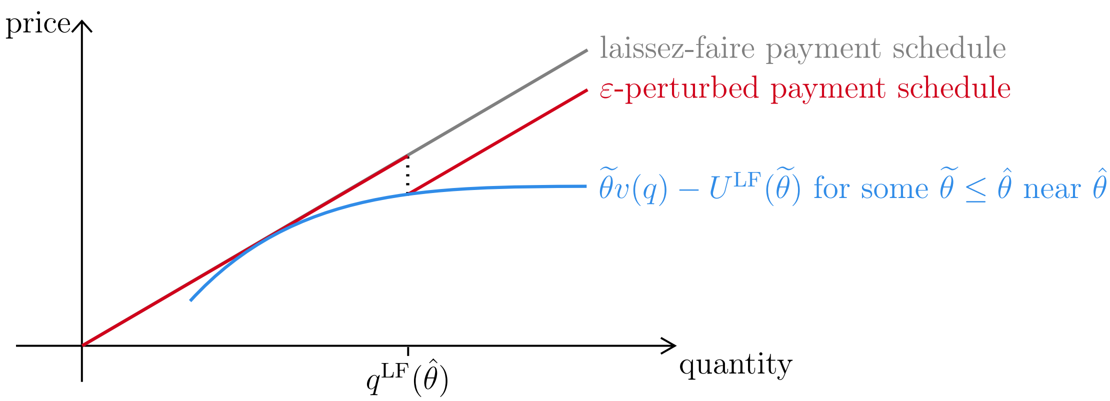

Figure 3: The ε-perturbed subsidy schedule, illustrating an indifference curve of type $\widetilde{\theta} \leq \hat{\theta}$ just indifferent between $q^{\mathrm{LF}}(\widetilde{\theta})$ and the subsidized $q^{\mathrm{LF}}(\hat{\theta})$. Types between $\widetilde{\theta}$ and $\hat{\theta}$ strictly prefer the subsidized $q^{\mathrm{LF}}(\hat{\theta})$.

3.2 Discussion

To understand the role of topping up in Theorem 1, we compare our result on the scope of intervention with similar results for two alternative models of private market interactions with government. In the first, which we call the *shutdown* benchmark, the social planner can shut down the private market for the good. In the second, called the *opt-in* benchmark, the social planner can require consumers to choose between participating in the subsidized market or the private market.

Our main conclusion is that the consumer's ability to top up in a private market generally limits the social planner's ability to redistribute. In the special case in which welfare weights are increasing in consumer demand, the topping up constraint is non-binding, and the optimal subsidy program with topping up coincides with the program enacted by a social planner who can require consumers to opt in or out of the subsidized market.

[12] We prove this formally in Appendix C.4.

Shutdown Benchmark. With the ability to shut down the private market for the good, the social planner can set prices higher than c for some units of the good to extract revenue from consumers with higher welfare weights and then use that revenue to subsidize consumers with lower welfare weights. As a result, 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 (see Appendix C.5 for details). This means that there is a wider scope of in-kind redistribution when the social planner can shut down the private market. For example, when ω is decreasing in θ and $\mathbf{E}[\omega] \leq \alpha$, a social planner unconstrained by a private market taxes all consumption, with higher taxes levied on higher consumption levels. By an argument similar to the one used in the previous subsection, 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, which is nonnegative because $\mathbf{E}[\omega] \leq \alpha$.

Opt-In Benchmark. In some redistribution programs (e.g., public housing and some public childcare and education programs), it is difficult for consumers to top up subsidized allocations in the private market. This restriction expands the social planner's set of feasible mechanisms by limiting consumers' ability to deviate: consumers need to decide whether to consume in the subsidized market or the private market.[13]

In a companion paper (Kang and Watt, 2024), we study the resulting subsidy design problem. In that case, the social planner offers a subsidy program if and only if the *maximum* welfare weight among all consumer types exceeds the opportunity cost of funds. Intuitively, 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.

When welfare weights are correlated with demand, the condition for subsidies without topping up matches the condition in Theorem 1 for markets with topping up. In fact, we show in our companion paper that the social planner never needs to enforce restrictions on topping up in order to implement the optimal mechanism—as a result, the optimal mechanisms we identify for increasing welfare weights are the same in both models. As we discuss in Section 4, this is because consumers with low types in the optimal mechanism effectively consume only in the private market,

[13] On the other hand, any feasible mechanism with topping up can be offered in the opt-in context by the social planner producing all units of the good and selling them at the subsidized payment schedule, which would lead all consumers to opt in (rather than the equivalent implementation of the subsidy mechanism with topping up, discussed in Section 2.4, in which the social planner pays subsidies for goods purchased in the private market).

while consumers with high types consume only in the subsidized market.

However, when welfare weights are negatively correlated with demand, the condition for offering subsidies with topping up is strictly more restrictive than in markets without topping up. This is because in the optimal mechanism with topping up and negative correlation, the social planner, who wants to target low-demand consumers, needs to also pay any subsidies offered to high-demand consumers, because they can always mimic the low-demand type and top up in the private market. If the social planner could require consumers to choose between participating in the subsidized market, the very high-demand types would not find it beneficial to mimic the low-demand type to receive the subsidy because it would require greatly reducing their consumption of the good, which cannot be topped up in the private market.

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:

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

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.[14] 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 (IR) 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 $J(\overline{\theta}) = \overline{\theta}$, reflecting the classic "no distortion at the top" property.

Ironing. We define the ironing operator $\overline{\phi} : [\underline{\theta}, \hat{\theta}] \to \mathbb{R}$ 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 $\Phi : [\underline{\theta}, \hat{\theta}] \to \mathbb{R}$ be defined by $\Phi(\theta) = \int_{\underline{\theta}}^{\theta} \phi(s) \, \mathrm{d}F(s)$, and let $\operatorname{co}\Phi$ be the *convex envelope* of Φ, which is the pointwise largest convex function satisfying for all $\theta \in [\underline{\theta}, \hat{\theta}]$, $\operatorname{co}\Phi(\theta) \leq \Phi(\theta)$. Then $\overline{\phi}$ is the monotone function satisfying

$$\text{for all } \theta \in [\underline{\theta}, \hat{\theta}], \qquad \int_{\underline{\theta}}^{\theta} \overline{\phi}(s) \, \mathrm{d}F(s) = \operatorname{co}\Phi(\theta).$$

We illustrate this ironing operation in Figure 4.

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

First, for any $s \in [\underline{\theta}, \theta]$, either $\overline{\phi}(s) = \phi(s)$ or s is contained in an *ironing interval*, on which $\overline{\phi}$ is constant and equal to that interval's F-weighted average of ϕ. This result is well-known but established formally in Lemma 2 in Appendix A.

Second, note that the definition of $\overline{\phi}$ depends on the domain of ϕ through the construction of $\operatorname{co}\Phi$, as illustrated in Figure 4. 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 Lemma 3 in Appendix A is that the value of an ironed function at θ 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 4.

In mechanism design, ironing is used to address violations of the monotonicity constraint. In

[14] 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.

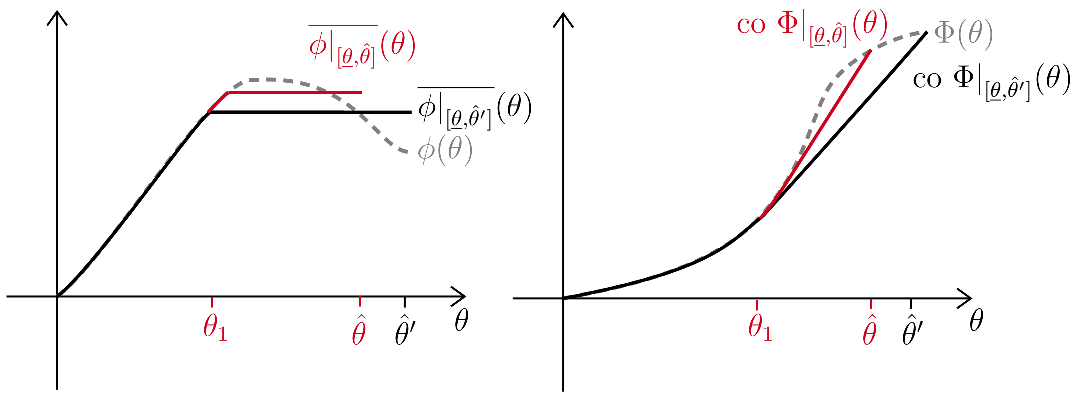

Figure 4: Illustrating the ironing operation applied to ϕ on domains $[\underline{\theta}, \hat{\theta}]$ (in red) and $[\underline{\theta}, \hat{\theta}']$ (in black). The ironed functions agree on $[\underline{\theta}, \theta_1]$, after which $\overline{\phi|_{[\underline{\theta},\hat{\theta}]}} > \overline{\phi|_{[\underline{\theta},\hat{\theta}']}}$.

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*

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

and $\kappa_+(\theta) = \inf \left\{ \widehat{\theta} \in \Theta : \widehat{\theta} \ge \theta, \text{ and } \widehat{\theta} \ge \overline{J|_{[\underline{\theta},\widehat{\theta}]}}(\widehat{\theta}) \right\}$ or $\overline{\theta}$ 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 5 illustrates the construction of the subsidy type $H(\theta)$ for two possible virtual welfare functions. 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 5(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 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 5(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 $\theta \geq \hat{\theta}$.

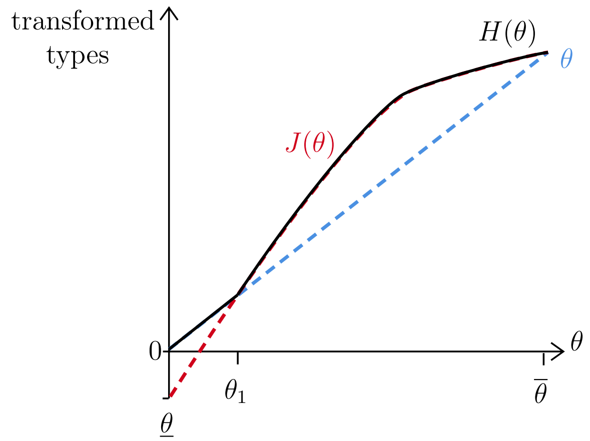

(a) In this case, $J(\theta)$ is increasing, and the subsidy type $H(\theta)$ is equal to the lower bound θ for $\theta \leq \theta_1$ and then $J(\theta)$ for $\theta \geq \theta_1$.

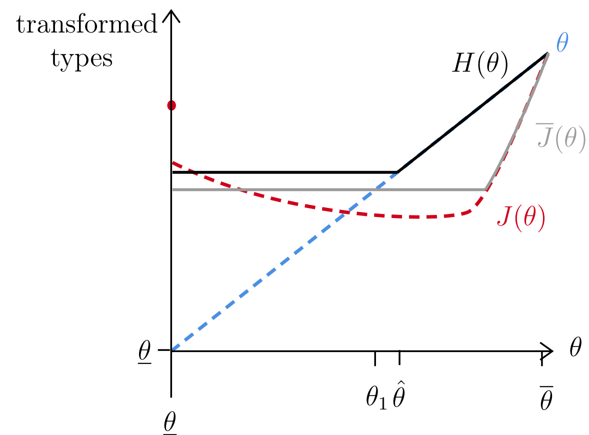

(b) In this case, $J(\theta)$ is nonmonotone with a point mass at $\underline{\theta}$, and the subsidy type $H(\theta)$ is equal to $\overline{J|_{[\underline{\theta},\hat{\theta}]}}$ for $\theta \leq \hat{\theta}$ and then θ for $\theta \geq \hat{\theta}$.

Figure 5: Constructing the subsidy type $H(\theta)$

Role of the Topping Up Constraint. To understand the role of 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. In Appendix C.5, we show that the optimal allocation rule in the shutdown problem is (cf. Toikka, 2011)

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

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}. This means that a social planner who can shut down the private market almost always finds it profitable to intervene in the market, reflecting our finding, discussed in Section 3.2, that the topping up constraint restricts the scope of redistribution by the social planner.

As is clear from Figure 5, the subsidy type in the problem with topping up is *not* obtained simply as the pointwise maximum of $\overline{J}(\theta)$ (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 $H(\theta) > \overline{J}(\theta)$. That distortion affects a larger set of agents than those for whom the lower-bound constraint would be binding at q^{SD}. In Figure 5(b), for example, $q^* > q^{\mathrm{SD}}$ for all consumer types, even though the (LB) constraint would be binding on q^{SD} only on $[\theta_1, \overline{\theta}]$. In other words, the optimal allocation is not obtained by relaxing (IC-T) and then enforcing the (LB) constraint 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 top up his allocation, 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 5, for example, if $q(\theta) = \max\{q^{\mathrm{LF}}(\theta), q^{\mathrm{SD}}(\theta)\}$, 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 $\overline{J}(\theta)$ 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.[15] Table 1 summarizes the optimal subsidies in each case.

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.

[15] 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.

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

Table 1: Dependence of the optimal mechanism on market primitives

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 markets with topping up.

Proposition 2 (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 2 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 because of the (IC-T) and (IR) constraints (otherwise, the consumers would prefer to consume only in the private market). 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 11 in Appendix A and offer an intuitive description here. Given the payments for the lowest type determined in Proposition 2, the allocation rule $q(\cdot)$ determines the slope of the payment schedule $T^{\text{S}}(z)$ 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 $T^{\text{S}}(z)$ 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 $T^{\text{S}}(z)$ 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 5(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 3 (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*

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

If $\min \omega \geq \alpha$*, then* $\theta_\alpha = \overline{\theta}$*, otherwise* $\theta_\alpha < \overline{\theta}$*. 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 3 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 6.

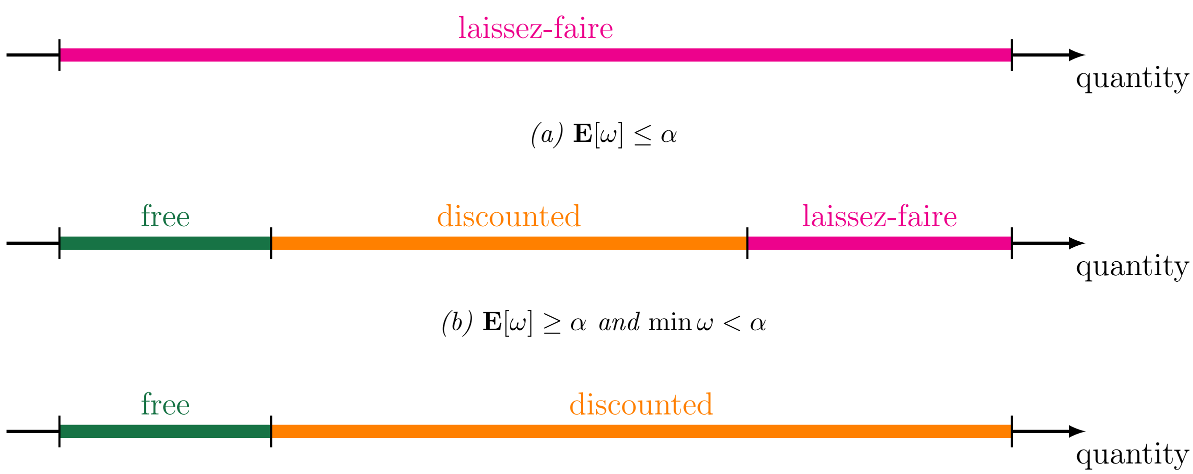

(a) $\mathbf{E}[\omega] \leq \alpha$

(b) $\mathbf{E}[\omega] \geq \alpha$ *and* $\min \omega < \alpha$

(c) $\mathbf{E}[\omega] \geq \alpha$ *and* $\min \omega \geq \alpha$

Figure 6: Marginal prices as a function of quantity for decreasing welfare weights.

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 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 5(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 4 (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 $\theta_\alpha = \overline{\theta}$ 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 $\theta_\alpha < \overline{\theta}$.

(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 4 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, $q^*(\underline{\theta})$, 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 4 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 7.

5.3 When Does The Planner Want To Restrict Topping Up?

As we discussed in Section 3.2, it may be possible for the social planner to restrict recipients in some subsidy programs from topping up 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

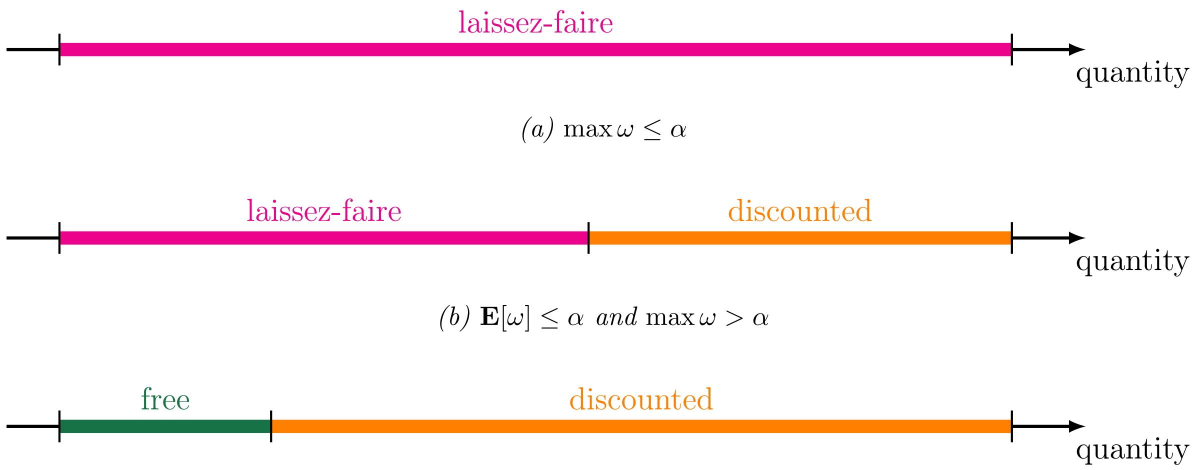

(a) $\max \omega \le \alpha$

(b) $\mathbf{E}[\omega] \le \alpha$ *and* $\max \omega > \alpha$

(c) $\mathbf{E}[\omega] > \alpha$

Figure 7: Marginal prices as a function of quantity when demand is positively correlated with need.

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-Kasparek, 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 top up to a higher-coverage "gold" plan (Healthcare.gov, 2024).

Restricting topping up reduces the set of constraints facing the social planner to (IC), (NLS) and (IR). Whereas the (LB) constraint in this paper implies that we need only enforce the (IR) 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 without topping up. Instead, we show in our companion paper (Kang and Watt, 2024) that the (IR) constraint takes the form of a majorization constraint on the allocation rule, requiring that

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

In our companion paper, we show that—in the case ω and θ are positively correlated—the optimal subsidy mechanism with no topping up satisfies the stricter topping up constraint (LB). As a result, there is no benefit to restricting topping up 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 positively correlated, the optimal subsidy mechanism with no topping up 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 no topping up to the mechanism identified in this paper. Intuitively, if the social planner can prohibit topping up, 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 without being able to top up in the private market.

5.4 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 5 (increasing preferences for redistribution)**.** *Suppose that there is a shift in the preference for redistribution from ω to $\widetilde{\omega}$ with $\widetilde{\omega}(\theta) \geq \omega(\theta)$ for all $\theta \in \Theta$, and/or a reduction in the opportunity cost of funds from α to $\widetilde{\alpha} < \alpha$. 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,*[16] *to* $\widetilde{H}(\theta) \geq H(\theta)$,

(b) *the set of subsidized consumers increases (in the sense of set inclusion),*

(c) *each consumer's total allocation increases, so* $\widetilde{q}^*(\theta) \geq q^*(\theta)$,

(d) *the total subsidy received by each consumer increases, and*

(e) *each consumer is better off overall, so* $\widetilde{U}^*(\theta) \geq U^*(\theta)$.

We prove Proposition 5 in Appendix C.6.

Intuitively, an increased preference for redistribution causes the virtual welfare of each type to increase to $\widetilde{J}(\theta) > J(\theta)$, resulting in a stronger preference for the social planner to distort consumption all types upwards. Proposition 5 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 8 (separately for the case in which ω is decreasing in θ and when it is increasing).

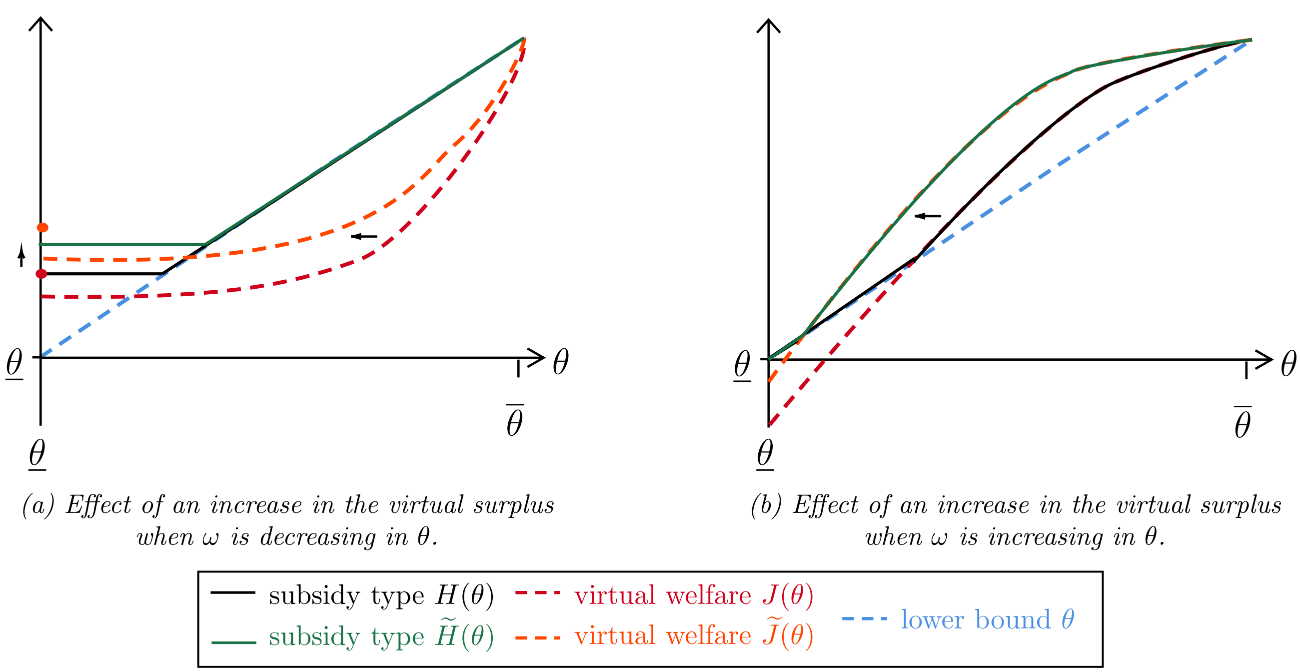

(a) Effect of an increase in the virtual surplus when ω is decreasing in θ.

(b) Effect of an increase in the virtual surplus when ω is increasing in θ.

Figure 8: Effects of a change in the virtual surplus on the subsidy type

[16] 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 $\widetilde{x}$ satisfies $\widetilde{x} \geq x$.

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 5 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 $\widetilde{\omega}$ are majorized by ω if for all $\hat{\theta} \in \Theta$, we have

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

and $\mathbf{E}[\widetilde{\omega}] = \mathbf{E}[\omega]$. In other words, the social planner places more weight on high-demand consumers under welfare weights $\widetilde{\omega}$ 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 9.

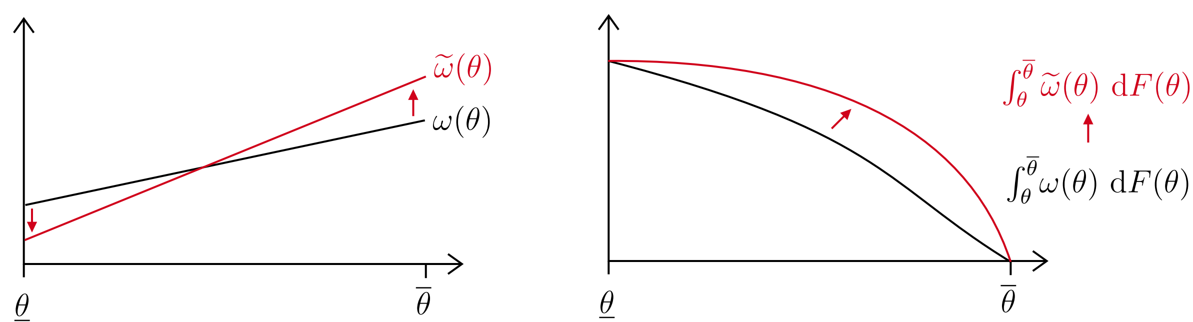

Figure 9: Increasing interdependence: $\widetilde{\omega}$ majorized ω, leading to a decrease in the distortion term.

A similar analysis as for Proposition 5 implies the following effect of increasing correlation between demand and need on the optimal subsidy mechanism.

Proposition 6 (effect of changing interdependence)**.** *Suppose that $\widetilde{\omega}$ is majorized by ω. Then, the total weighted surplus is greater and the optimal subsidy mechanism is more generous given $\widetilde{\omega}$ than ω, in that:*

(a) each consumer's subsidy type increases, to $\widetilde{H}(\theta) \geq H(\theta)$*,*

(b) the set of subsidized consumers is larger (in the sense of set inclusion),

(c) each consumer's total allocation increases, so $\widetilde{q}^*(\theta) \leq q^*(\theta)$*,*

(d) the total subsidy received by each consumer increases, and

(e) each consumer's utility increases, so $\widetilde{U}^*(\theta) \leq U^*(\theta)$*.*

We prove Proposition 6 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 6 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 $\widetilde{c} < c$, say, or by a change in the valuation function v to a new valuation function $\widetilde{v}$ with $\widetilde{v}\,'(q) \geq v'(q)$ 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 7 (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 7 in Appendix C.6.

Holding all else equal, Proposition 7 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 5.4, 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 reduces the extent of topping up by consumers (which slackens the topping up constraint for the social planner).

To incorporate these effects, we now suppose that there is a strictly convex and differentiable industry cost function $C : [0, A] \to \mathbb{R}$, where $C(Q)$ is the total cost to all suppliers of producing $Q = \int_{\Theta} q(\theta)\, \mathrm{d}F(\theta)$ units of the good. We maintain the assumption that the industry is competitive, so that this cost function leads to an industry supply curve $S(p) = \max_{Q\in[0,A]} [pQ - C(Q)]$, 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).

As discussed above, a changing price in the private market affects each consumer's ability to top up his consumption of the good. We capture this possibility in revised incentive compatibility constraints and individual rationality constraints. For any mechanism (q, t), let $p = S^{-1}\left(\int_\Theta q(\theta)\ \mathrm{d}F(\theta)\right)$ be the equilibrium price of the good induced by the mechanism, then the revised private market constraint is

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

and the revised individual rationality constraint is

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

Proposition 8 (optimal subsidy with equilibrium effects)**.** *The allocation rule solving*

$$\max_{(q,t)} \int_\Theta \Big[\omega(\theta)\left[\theta v(q(\theta)) - t(\theta)\right] + \alpha t(\theta)\Big]\ \mathrm{d}F(\theta) - \alpha C\left(\int_\Theta q(\theta)\ \mathrm{d}F(\theta)\right),$$
$$\textit{such that } (q, t) \textit{ satisfies } (\mathrm{IC}),\ (\mathrm{IR'}),\ (\mathrm{NLS}) \textit{ and } (\mathrm{PM'}),$$

is characterized by

$$q^*(\theta) = D\left(p, H(\theta)\right),\ \textit{ and } p = S^{-1}\left(\int_\Theta q^*(\theta)\ \mathrm{d}F(\theta)\right),$$

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

The proof of Proposition 8 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 8 implies that the optimal allocation rule can be found by constructing a *subsidized demand* curve

$$D(p) = \int_\Theta D\left(p, H(\theta)\right)\ \mathrm{d}F(\theta)$$

and identifying its intersection p^* with the industry supply curve, determining the competitive market price, as illustrated in Figure 10. 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 $\mathbf{E}[\omega] \geq \alpha$).

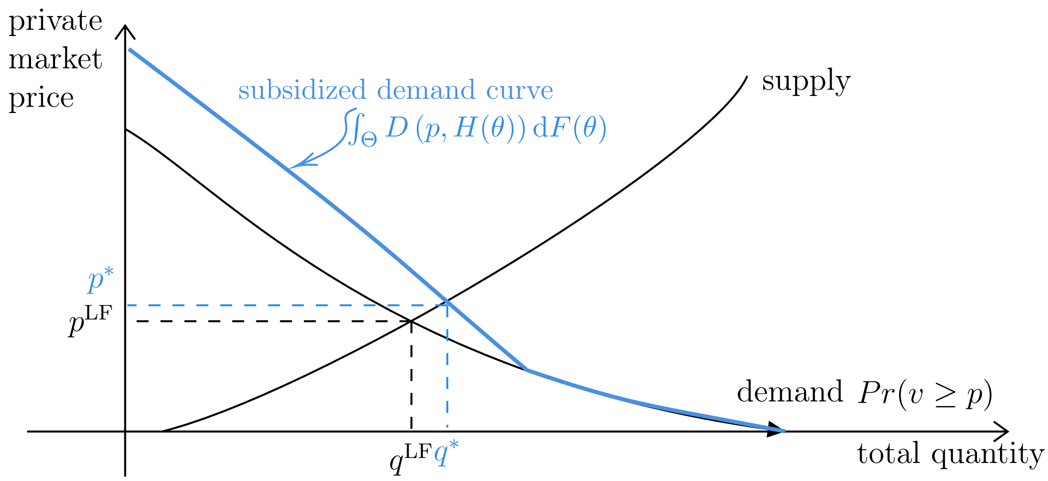

Figure 10: Illustrating the Subsidized Demand Curve with Equilibrium Effects

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 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 topping up. 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,[17] so the subsidy design problem takes the form

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

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

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

[17] Otherwise, if the social planner must 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 $c(1+\tau)$.

and a lower-bound constraint incorporating the tax τ

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

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 9 (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 $H_\tau(\theta) = \overline{H_\tau^{\mathrm{T}}}(\theta)$ *and*

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

We prove Proposition 9 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

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

subject to (IC), (NLS), (IR_t) and (LB_t). In that case, the optimal subsidy allocation is determined as in Proposition 9, 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

$$\max_{(q,t)} \int_\Theta [\omega(\theta)v(q(\theta)) - t(\theta)] \; \mathrm{d}F(\theta),$$

subject to (IC-T), (IR), (NLS), and the additional constraint

$$\int_\Theta [cq(\theta) - t(\theta)] \; \mathrm{d}F(\theta) \leq B,$$

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

A microfoundation for the welfare weights $\omega(\theta)$ was offered by Pai and Strack (2024) as the expected marginal value for money of a consumer with concave preferences over money. 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 $\varphi : \mathbb{R} \to \mathbb{R}$. Suppose that the social planner solves

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

subject to (IC-T), (IR), 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 (SUB) 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.[18] 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

[18] 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).

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 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.[19]

Applying our results to this context,[20] 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.[21] These fare capping programs are consistent with the volume discounts in the optimal subsidy mechanism in the positive correlation case, identified in Section 5.[22]

[19] 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.

[20] 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.

[21] 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).

[22] We note that zero marginal prices in the payment schedule (as in fare capping) arise in the limit as $q \to \infty$ of

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*, 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.[23] 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

the optimal subsidy schedule when F has decreasing hazard rate and ω is increasing in θ.

[23] 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).

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 the program's history.[24] 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.

[24] 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.

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 those private markets to top up their consumption of subsidized goods limits the redistributive power of the subsidy program. When consumers can top up, the social planner must account for the fact that subsidies must increase in demand for the good, which makes it easier to intervene when subsidy designers aim 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 that 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. ⓡ E. Budish ⓡ P. Dworczak ⓡ S. D. Kominers (2024a): "An economic framework for vaccine prioritization," *The Quarterly Journal of Economics*, 139, 359–417.

Akbarpour, M. ⓡ P. Dworczak ⓡ 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.

Arrow, K. J. (1951): "An Extension of the Basic Theorems of Classical Welfare Economics," in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability* ed. by Neyman, J. Berkeley, CA: University of California Press, 507–532.

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., and J. A. Mirrlees (1972): "Optimal taxation and public production: I," *The American Economic Review*, 62, 238–238.

Downs, A. (1957): *An Economic Theory of Democracy*: Harper & Row.

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

Dworczak, P. ⓡ S. D. Kominers ⓡ M. Akbarpour (2021): "Redistribution Through Markets," *Econometrica*, 89, 1665–1698.

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.

Helsenorge (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.

Maskin, E., and J. Riley (1984): "Monopoly with Incomplete Information," *RAND Journal of Economics*, 15, 171–196.

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.

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.

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.

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. (1982): "Self-selection and Pareto efficient taxation," *Journal of public economics*, 17, 213–240.

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.

Woody, J., Z. Zhao, R. Lund, and T.-L. Wu (2024): "A forensic statistical analysis of fraud in the federal food stamp program," *The Annals of Applied Statistics*, 18, 2486–2510.

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 reformulate the (IC-T) constraint to obtain a convex program and then discuss the generalization 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 Reformulating the Subsidy Design Problem as a Convex Program

Proposition 10 (reformulating the subsidy design problem)**.** *The subsidy design problem* (SUB) *is equivalent to the following convex program:*

$$\max_{q:\Theta\to[0,A]} \ \alpha \int_{\Theta} [J(\theta)v(q(\theta)) - cq(\theta)] \ \mathrm{d}F(\theta) + (\textit{terms independent of } q), \tag{OPT}$$
$$\textit{such that } q \textit{ is nondecreasing and satisfies for all } \theta \in \Theta,\ q(\theta) \geq q^{\mathrm{LF}}(\theta),$$

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

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

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

Proof. Using Proposition 1, we have that the subsidy design problem takes the form

$$\max_{(q,t)} \int_{\Theta} \Big\{\omega(\theta) \underbrace{[\theta v(q(\theta)) - t(\theta)]}_{\text{consumer surplus}} - \alpha \underbrace{[cq(\theta) - t(\theta)]}_{\text{total costs}} \Big\} \ \mathrm{d}F(\theta), \tag{SUB}$$
$$\text{such that } (q,t) \text{ satisfies (IR), (IC), (LB) and (NLS).}$$

We apply standard mechanism design techniques to eliminate various constraints and express (SUB) 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 (IR) constraint only for the lowest type, that is, to ensure that $\underline{U} := \underline{\theta}v(q(\underline{\theta})) - t(\underline{\theta}) \geq U^{\mathrm{LF}}(\underline{\theta})$. To see that, note that

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

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

so $U(\theta) \geq U^{\mathrm{LF}}(\theta)$ 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

$$\max_{(q,t)} \left\{ (\mathbf{E}[\omega] - \alpha)\, \underline{U} + \alpha \int_{\Theta} \left[\left(\theta + \frac{\int_{\theta}^{\overline{\theta}} [\omega(s) - \alpha]\ \mathrm{d}F(s)}{\alpha f(\theta)} \right) v(q(\theta)) - cq(\theta) \right] \mathrm{d}F(\theta) \right\},$$
$$\text{such that } q \in \mathcal{Q},\ \underline{U} \geq U^{\mathrm{LF}}(\underline{\theta}),\ \text{and } (q,t) \text{ satisfies (NLS) and (LB).}$$

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 $\theta, \theta' \in \Theta$,

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

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) constraint binds and $\underline{U} = U^{\mathrm{LF}}(\underline{\theta})$. 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 $\mathbf{E}[\omega] > \alpha$:

$$\max_{q \in \mathcal{Q}\colon q \geq q^{\mathrm{LF}}} \alpha \int_{\Theta} [J(\theta)v(q(\theta)) - cq(\theta)] \; \mathrm{d}F(\theta), \tag{OPT-L}$$

where

$$J(\theta) = \theta + \frac{\int_{\theta}^{\overline{\theta}} [\omega(s) - \alpha] \; \mathrm{d}F(s) + (\mathbf{E}[\omega(\theta)] - \alpha)\underline{\theta}\delta_{\underline{\theta}}}{\alpha f(\theta)}.$$

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

$$\max_{q \in \mathcal{Q}\colon q \geq q^{\mathrm{LF}}} \alpha \int_{\Theta} [J(\theta)v(q(\theta)) - cq(\theta)] \; \mathrm{d}F(\theta) + (\mathbf{E}[\omega] - \alpha)\, \underline{\theta} v(q^{\mathrm{LF}}(\underline{\theta})), \tag{OPT-H}$$

where

$$J(\theta) = \theta + \frac{\int_{\theta}^{\overline{\theta}} [\omega(s) - \alpha] \; \mathrm{d}F(s)}{\alpha f(\theta)}.$$

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 10. □

A.2 Generalizing (OPT) and Theorem 2

In this subsection, we introduce a generalization of the (LB) constraint that requires q to exceed pointwise an arbitrary continuous $q^{\mathrm{L}} \in \mathcal{Q}$.[25] 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 $\lim_{z\to\infty} v'(z) = 0$ and $\lim_{z\to 0} v'(z) = +\infty$.

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

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

In this section, we write $D(\theta)$ for the demand of type θ (suppressing the price c, which is held

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

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 $L(\theta) = D^{-1}(q^{\mathrm{L}}(\theta))$, which is the type demanding $q^{\mathrm{L}}(\theta)$ 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)*,* $q^* : \Theta \to [0, A]$*, 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

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

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

Note that Theorem 2 follows as its consequence because if $q^{\mathrm{L}}(\theta) = q^{\mathrm{LF}}(\theta)$, 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.3 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

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

Recall that on any ironing interval $[\theta_1, \theta_2]$ of ϕ, for all $\theta \in [\theta_1, \theta_2]$, we have $\overline{\phi}(\theta) = \mu(\theta_1, \theta_2)$. 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 1 (Interrupted Ironing). *Let $\overline{\phi}(\theta) = c$ 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 $\int_{\theta_1^*}^{\theta_2} \phi(s)\ \mathrm{d}F(s) \geq \int_{\theta_1^*}^{\theta_2} c\ \mathrm{d}F(s)$, so $\mu(\theta_1^*, \theta_2) \geq c$. Similarly, since $\int_{\theta_1^*}^{\theta_1} \phi(s)\ \mathrm{d}F(s) \geq \int_{\theta_1^*}^{\theta_1} c\ \mathrm{d}F(s)$ and $\int_{\theta_1^*}^{\theta_2^*} \phi(s)\ \mathrm{d}F(s) = \int_{\theta_1^*}^{\theta_2^*} c\ \mathrm{d}F(s)$, we have $\int_{\theta_1}^{\theta_2^*} \phi(s)\ \mathrm{d}F(s) \leq \int_{\theta_1}^{\theta_2^*} c\ \mathrm{d}F(s)$, 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 2 (max-min formula for ironing). *For any $\phi : \Theta \to \mathbb{R}$ 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 $\overline{\phi}(\theta) = c$ 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 $\overline{\phi}$, we have $\overline{\phi}(\theta_2^*) = c' > c$. If θ_2^* is inside an ironing interval of ϕ, Lemma 1 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 $\overline{\phi}$, 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 1, 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 $\overline{\phi}$ (as the slope of the convex envelope of $\operatorname{co} \Phi$), the interval $[\theta_1^*, \theta_2^*]$ is an ironing interval of $\overline{\phi}$ containing θ as well. $\square$

The expression for the ironing operator as the solution of a minimax problem obtained in Lemma 2 allows us to establish the monotonicity of the ironing operator in the domain of its argument:

Lemma 3 (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 3 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 2, 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.4 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).

Lemma 4 (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)

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

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

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

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

We introduce additional notation for the proof of Lemma 4. 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 $\Theta^1, \Theta^2, \ldots$ so $\Theta = \cup_i \Theta^i$ and let $\operatorname{int} \Theta^i = (\theta_-^i, \theta_+^i)$.

Let

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

matching the expression in part (a) of Lemma 4. To establish the *equivalence* of our expressions for H, we need to show that $\widetilde{H} = H = \overline{H^T}$. Once we have proven those equivalent representations, to establish *feasibility* of q^*, we need only establish that at least one of $\widetilde{H}$, 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 $\widetilde{H} = H$. This is clear for any θ in a type (a) interval, so we need only show for θ in type (b) intervals that $\widetilde{H}(\theta) = \overline{J|_{[\theta_-^i,\theta_+^i]}}(\theta) = \overline{J|_{[\underline{\theta},\theta_+^i]}}(\theta) = H(\theta)$. By Lemma 3, it suffices to show that θ_-^i is not contained in an ironing interval in $[\underline{\theta}, \theta_+^i]$ as this will imply $\overline{J|_{[\theta_-^i,\theta_+^i]}}(\theta) = \overline{J|_{[\underline{\theta},\theta_+^i]}}(\theta)$, as required. Suppose otherwise, then there must exist an ironing interval of $J|_{[\underline{\theta},\theta_+^i]}$ containing θ_-^i ending at some $\hat{\theta} \in \Theta^i$. Because $\hat{\theta} \in \Theta^i$, we must have that $\overline{J|_{[\underline{\theta},\theta_+^i]}}(\hat{\theta}) \geq L(\hat{\theta})$ which implies $\overline{J|_{[\underline{\theta},\theta_+^i]}}(\theta_-^i) \geq L(\hat{\theta}) \geq L(\theta_-^i)$. But then for $\theta \leq \theta_-^i$ contained in the same ironing interval, we must have (by Lemma 3) that $\overline{J|_{[\underline{\theta},\theta]}}(\theta) \geq \overline{J|_{[\underline{\theta},\theta_+^i]}}(\theta) \geq L(\hat{\theta}) \geq L(\theta)$, which contradicts the fact that θ is in a type (a) interval.

We now show that $\widetilde{H}$ 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 $\widetilde{H}$ match at the meeting points of adjacent intervals. But by construction, for any type (b) interval $[\theta_-^i, \theta_+^i]$, we have $\overline{J|_{[\underline{\theta},\theta_-^i]}}(\theta_-^i) = L(\theta_-^i)$ and $\overline{J|_{[\underline{\theta},\theta_+^i]}}(\theta_+^i) = L(\theta_+^i)$, and then, by our previous finding, we have that $\widetilde{H}(\theta_-^i) = \overline{J|_{[\theta_-^i,\theta_+^i]}}(\theta_-^i) = L(\theta_-^i)$, and $\widetilde{H}(\theta_+^i) = \overline{J|_{[\theta_-^i,\theta_+^i]}}(\theta_+^i) = L(\theta_+^i)$, as required.

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

To obtain that result, we first show that $\widetilde{\mathcal{H}}$ is a minorant of $\mathcal{H}^{\mathrm{T}}$. To see this, note that within any interval Θ^i, the function $\theta \mapsto \int_{\theta_-^i}^{\theta} \widetilde{H}(s)\, \mathrm{d}F(s)$ is the convex envelope of $\theta \mapsto \int_{\theta_-^i}^{\theta} H^{\mathrm{T}}(s)\, \mathrm{d}F(s)$. This means that for all $\theta \in \Theta^i$, $\int_{\theta_-^i}^{\theta} H(s)\, \mathrm{d}F(s) \leq \int_{\theta_-^i}^{\theta} H^{\mathrm{T}}(s)\, \mathrm{d}F(s)$ with equality at $\theta = \theta_+^i$. But then, applying those facts from left to right, we obtain that $\widetilde{\mathcal{H}}(\theta) \leq \mathcal{H}^{\mathrm{T}}(\theta)$ for all $\theta \in \Theta$, and that $\widetilde{\mathcal{H}}(\theta) = \mathcal{H}^{\mathrm{T}}(\theta)$ for any θ at the endpoint of an interval. That implies that $\widetilde{\mathcal{H}}$ is a minorant of $\mathcal{H}^{\mathrm{T}}$ 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 $\mathcal{H}^{\mathrm{T}} \geq \mathcal{H}' \geq \mathcal{H}$ with $\mathcal{H}' \neq \mathcal{H}$). Since $\widetilde{\mathcal{H}} = \mathcal{H}^{\mathrm{T}}$ at the endpoint of each interval Θ^i, $\mathcal{H}' = \mathcal{H}^{\mathrm{T}}$ there as well. So if $\mathcal{H}'(\theta) > \widetilde{\mathcal{H}}(\theta)$ for some $\theta \in \Theta$, it must be for some $\theta \in \mathrm{int}(\Theta^i)$. But then $\mathcal{H}'|_{\Theta^i}$ would be a convex minorant of $\mathcal{H}^{\mathrm{T}}|_{\Theta^i}$ with $\mathcal{H}'|_{\Theta^i} > \widetilde{\mathcal{H}}|_{\Theta^i}$, which would contradict the fact that $\widetilde{\mathcal{H}}|_{\Theta^i}$ is the convex envelope of $\mathcal{H}^{\mathrm{T}}|_{\Theta^i}$. This implies that $\widetilde{\mathcal{H}}$ is the convex envelope of $\mathcal{H}^{\mathrm{T}}$, and so $\widetilde{H} = H$.

Finally, having shown that $\widetilde{H}(= H)$ is nondecreasing, for the feasibility of $q^* = D(c, H(\theta))$, we need only show that for all $\theta \in \Theta^i$, $\widetilde{H}(\theta) \geq L(\theta)$. 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 $\widetilde{H}(\theta) \geq L(\theta)$ as well. □

A.5 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) 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

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

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). The optimality condition of (OPT) may be written as a variational inequality: for any other feasible q, we have that

$$\int_\Theta [J(\theta)v(q^*(\theta)) - cq^*(\theta)]\ \mathrm{d}F(\theta) \geq \int_\Theta [J(\theta)v(q(\theta)) - cq(\theta)]\ \mathrm{d}F(\theta). \tag{VI}$$

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

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

it suffices to sign the difference in these inequalities:

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

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

$$\int_\Theta \left[J(\theta) - H^{\mathrm{T}}(\theta)\right] [v(q^*(\theta)) - v(q(\theta))]\ \mathrm{d}F(\theta) + \int_\Theta \left[H^{\mathrm{T}}(\theta) - H(\theta)\right] [v(q^*(\theta)) - v(q(\theta))]\ \mathrm{d}F(\theta) \geq 0,$$

and sign the terms in that sum individually.

First, we establish that

$$\int_\Theta [H^{\mathrm{T}}(\theta) - H(\theta)]v(q^*(\theta))\ \mathrm{d}F(\theta) = 0.$$

Intuitively, that follows because $H = \overline{H^{\mathrm{T}}}$, and on any interval where $H \neq H^{\mathrm{T}}$, 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 $c \in \operatorname{im} H$, and let $\Theta_c := \{\theta \in \Theta : H(\theta) = c\}$. Because $H = \overline{H^{\mathrm{T}}}$, either $c = H^{\mathrm{T}}(\theta) = H(\theta)$ for all $\theta \in \Theta_c$, or $\int_{\Theta_c} [H^{\mathrm{T}}(\theta) - H(\theta)]\ \mathrm{d}F(\theta) = 0$. In either case, since $v(q^*(\theta))$ is constant on Θ_c, we have

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

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

Second, we show that for any feasible q

$$\int_\Theta [H(\theta) - H^{\mathrm{T}}(\theta)]v(q(\theta))\ \mathrm{d}F(\theta) \geq 0.$$

Intuitively, wherever $H \neq H^{\mathrm{T}}$, 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

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

which is nonnegative because $v(q(\theta))$ is increasing in θ for any feasible q and $\int_{\underline{\theta}}^{\theta} H^{\mathrm{T}}(s)\ \mathrm{d}F(s) \geq \int_{\underline{\theta}}^{\theta} H(s)\ \mathrm{d}F(s)$ for all $\theta \in \Theta$ (with equality at $\underline{\theta}$ and $\overline{\theta}$) 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

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

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 $H^{\mathrm{T}}(\theta) = L(\theta)$. But by Lemma 3, $\overline{J|_{[\underline{\theta},\theta]}}(\theta) \geq J(\theta)$, so $L(\theta) \geq J(\theta)$, and thus $[J(\theta) - H^{\mathrm{T}}(\theta)]$ is nonpositive. For the second term in the product, where $H^{\mathrm{T}}(\theta) = L(\theta)$, we have via Lemma 4 that $H(\theta) = L(\theta)$, so $q^*(\theta) = q^{\mathrm{L}}(\theta)$. As a result, $q^*(\theta) = q^{\mathrm{L}}(\theta) \leq q(\theta)$ 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.

□

A.6 Proof of Proposition 2 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 2 which pins down the payment made by the lowest type $\underline{\theta}$.

Proof. As shown in Appendix A.1 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) constraint binds for $\underline{\theta}$. Because $J(\underline{\theta}) < \theta$, we have $H(\underline{\theta}) = \underline{\theta}$, so $q^*(\underline{\theta}) = D(c, \underline{\theta}) = q^{\mathrm{LF}}(\underline{\theta})$. 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 $q^*(\underline{\theta}) = q^{\mathrm{LF}}(\underline{\theta})$. As a result, the social planner is indifferent between feasible payments for $q^{\mathrm{LF}}(\underline{\theta})$ units of the good, which by the (NLS) and (IR) constraints imply $0 \leq t(\underline{\theta}) \leq c q^{\mathrm{LF}}(\underline{\theta})$. □

Proposition 11 (payment schedule)**.** *The following subsidized payment schedule implements the optimal subsidy allocation for quantities $z \in \operatorname{im} q^*$:*

$$T^{\mathrm{S}}(z) = \begin{cases} \displaystyle\int_{q^*(\underline{\theta})}^{z} (\theta^*)^{-1}(\hat{z}) v'(\hat{z}) \, \mathrm{d}\hat{z} & \text{if } \mathbf{E}[\omega] > \alpha, \\ \displaystyle\gamma + \int_{q^{\mathrm{LF}}(\underline{\theta})}^{z} (\theta^*)^{-1}(\hat{z}) v'(\hat{z}) \, \mathrm{d}\hat{z} & \text{if } \mathbf{E}[\omega] = \alpha, \\ \displaystyle c q^{\mathrm{LF}}(\underline{\theta}) + \int_{q^{\mathrm{LF}}(\underline{\theta})}^{z} (\theta^*)^{-1}(\hat{z}) v'(\hat{z}) \, \mathrm{d}\hat{z} & \text{if } \mathbf{E}[\omega] < \alpha, \end{cases}$$

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

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

Proof. By the Milgrom and Segal (2002) envelope theorem, we have

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

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

$$\frac{d}{dz} T^{\mathrm{S}}(z) = \frac{\mathrm{d}t^*/\mathrm{d}\theta}{\mathrm{d}q^*/\mathrm{d}\theta} = (\theta^*)^{-1}(z) v'(z),$$

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

Note that Proposition 11 implies that the total subsidy rule satisfies

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}z} S(z) &= c - \frac{\mathrm{d}}{\mathrm{d}z} T^{\mathrm{S}}(z) \\
&= c - (\theta^*)^{-1}(z) v'(z) \\
&= c - (\theta^*)^{-1}(z) \frac{c}{H((\theta^*)^{-1}(z))},
\end{aligned}$$

which is nonnegative because $H(\theta) \geq \theta$, according with our finding in Proposition 1 that total subsidy payments are increasing in quantity.

B Structure of Optimal Subsidies

In this section, we derive the structure of the optimal subsidy mechanism for increasing and decreasing ω, proving Proposition 3 and Proposition 4 stated in Section 5.2.

B.1 Proof of Proposition 3 (Negative Correlation)

To derive Proposition 3, 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 $\frac{\int_\theta^{\overline{\theta}}[\omega(s)-\alpha]\,\mathrm{d}F(s)}{\alpha f(\theta)}$ for $\theta \in (\underline{\theta}, \overline{\theta}]$.

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 $\int_\theta^{\overline{\theta}}[\omega(s)-\alpha]\,\mathrm{d}F(s)$ is quasiconvex in θ, while it is zero at $\overline{\theta}$ 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 11.[26]

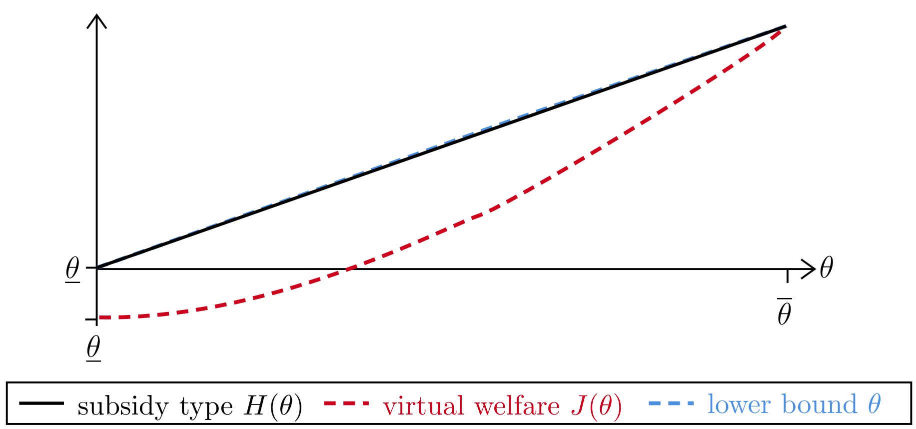

Figure 11: Subsidy types for decreasing welfare weights and a high opportunity cost of funds. In this case, the lower-bound constraint binds for all types.

Low opportunity cost of funds: $\mathbf{E}[\omega] > \alpha$ Proposition 2 implies that there is always public provision of the good when $\mathbf{E}[\omega] > \alpha$. To determine the quantity of public provision and any

[26] In Figures 11, 12, 13, and 14, 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 $\int_{\theta}^{\overline{\theta}}[\omega(s)-\alpha]\,\mathrm{d}F(s)$ is quasiconvex in θ, while it is zero at $\overline{\theta}$ 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 12(a), so $q^* = q^{\mathrm{SD}}$.

(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 $[\hat{\theta}, \overline{\theta}]$, as shown in Figure 12(b)(i). Writing $\hat{\theta} = \theta_\alpha$ as in Table 1, $q^*(\theta) = q^{\mathrm{SD}}(\theta)$ for $\theta \leq \theta_\alpha$, and $q^*(\theta) = q^{\mathrm{LF}}(\theta)$ for $\theta \geq \theta_\alpha$, with the associated payment schedule derived using Proposition 11, 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) $\overline{J|_{[\underline{\theta},\hat{\theta}]}}(\hat{\theta}) > J(\theta) = \hat{\theta}$ (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 $[\theta_\alpha, \overline{\theta}]$, as in Figure 12(b)(ii). This implies $q^* \geq q^{\mathrm{SD}}$, with strict inequality on the ironing interval of $J|_{[\underline{\theta},\theta_\alpha]}$ ending at θ_α, and $q^* = q^{\mathrm{LF}}$ on $[\theta_\alpha, \overline{\theta}]$. In this case, the payment schedule, derived using Proposition 11, involves a free allocation of $q^{\mathrm{LF}}(\theta_\alpha)$ units for all agents and laissez-faire prices for additional consumption beyond that amount.

B.2 Proof of Proposition 4 (Positive Correlation)

We again derive Proposition 4 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 $\int_{\theta}^{\overline{\theta}}[\omega(s)-\alpha]\,\mathrm{d}F(s)$ is quasiconcave in θ, while it is zero at $\overline{\theta}$ and nonpositive at $\underline{\theta}$ because $\mathbf{E}[\omega] \leq \alpha$. This leads to two possibilities for the sign of the distortion term:

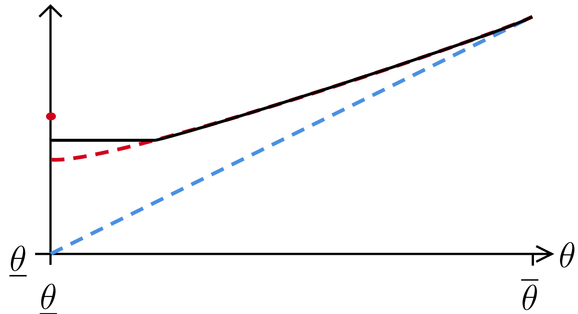

(a) Here $\omega(\theta) \geq \alpha$ for each θ, so the social planner distorts each type's allocation upward, with a free endowment for all types and quantity-dependent subsidies for additional units of the good.

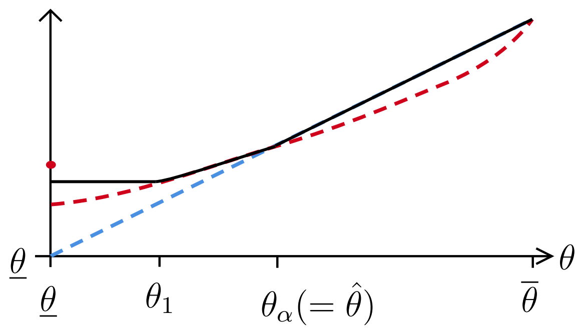

(b)(i) Here, the social planner distorts lower types' consumption upward, via a free endowment of $q^{\mathrm{LF}}(J(\theta_1))$ and limited quantity-dependent subsidies for additional units. Consumption of types $\theta \geq \theta_\alpha$ is undistorted.

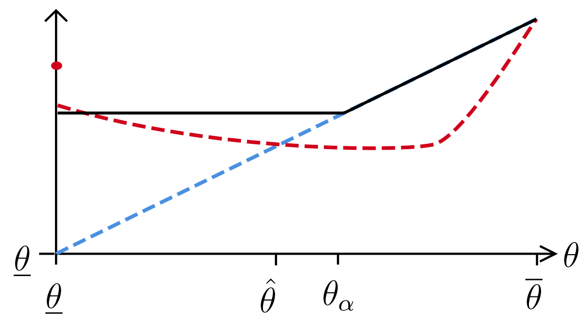

(b)(ii) In this example, the social planner distorts lower types' consumption upward, via a free endowment of $q^{\mathrm{LF}}(\theta_\alpha)$, which can strictly exceed $q^{\mathrm{SD}}(\theta)$ for $\theta \leq \theta_\alpha$. Consumption of types $\theta \geq \theta_\alpha$ is undistorted.

Figure 12: 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 13(a), so $q^* = q^{\mathrm{LF}}$, 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 $\theta \in [\theta_\alpha, \overline{\theta}]$. This means $H(\theta) = \overline{J}(\theta)$ for $\theta \geq \theta_\alpha$, and $H(\theta) = \theta$ for $\theta \leq \theta_\alpha$, as illustrated in Figure 13(b). As a result, $q^*(\theta) = q^{\mathrm{LF}}(\theta)$ for $\theta \leq \theta_\alpha$, and $q^*(\theta) = q^{\mathrm{SD}}(\theta)$ for $\theta \geq \theta_\alpha$. By Proposition 11, the social planner implements this allocation using subsidies for consumption above a certain minimum level, as illustrated in Figure 1(b).

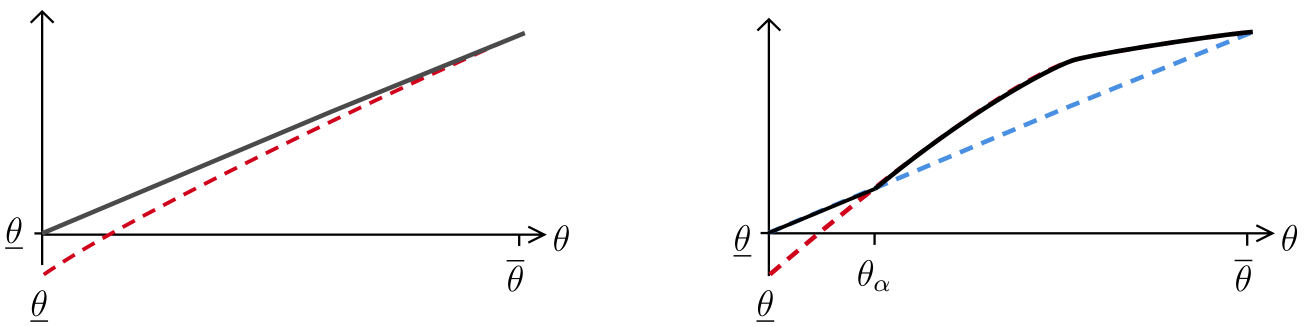

(a) In this case, the social planner wants to distort consumption for all types downwards, but is prevented from doing so by the (LB) constraint, so $q^* = q^{\mathrm{LF}}$.

(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 13: 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 2 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 $\int_\theta^{\overline{\theta}} [\omega(s) - \alpha] \, \mathrm{d}F(s)$ is quasiconcave in θ, while it is zero at $\overline{\theta}$ 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 14, the subsidy type equals the ironed virtual welfare, so $q^* = q^{\mathrm{SD}}$.

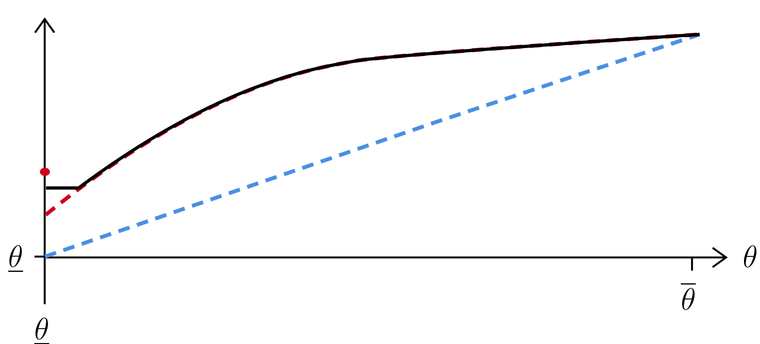

subsidy type $H(\theta)$ virtual welfare $J(\theta)$ lower bound θ

Figure 14: 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 $M : \Theta \to \Delta([0, A] \times \mathbb{R})$ 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,
$$\mathbf{E}_{(q,t)\sim M(\theta)}[\theta v(q) - t] \geq U^{\mathrm{LF}}(\theta),$$
- 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
$$\Pr_{(q,t)\sim M(\theta)}[q \geq q^{\mathrm{LF}}(\theta)] = 1,$$
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

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

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

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

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 $q : \mathbb{R} \to [0, A]$ 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 $\widetilde{q}$ defined by $v \circ \widetilde{q}(\theta) = \frac{1}{2}\left[v(q_1(\theta)) + v(q_2(\theta))\right]$. Clearly, $\widetilde{q}$ 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 $\widetilde{q}$, contradicting the optimality of q_1 and q_2.

C.3 Proof of Proposition 1 (Equivalent Representations of (IC-T))

Proof. We first show that *(i)* implies *(ii)*. Suppose otherwise: that (q, t) satisfied the (IC-T) constraint but could not be implemented by a subsidized payment schedule $T^{\mathrm{S}}(z)$ with $S(q) = cq - T^{s}(q)$ positive and nondecreasing. In that case, because as discussed in Section 2.4, any mechanism can be implemented using a nondecreasing payment outside of $\operatorname{im} q$, it must be the case that $t(\theta') < t(\theta)$ for some $\theta' \neq \theta$ with $q(\theta') > q(\theta)$. But then a consumer of type θ would be better off reporting type θ', because doing so would result in a higher consumption level $q(\theta')$ at a lower cost $t(\theta')$, which can only increase his utility. But that violates the (IC-T) constraint.

For the converse, suppose that T^{S} implements the mechanism (q, t) and $S(z) = cz - T^{\mathrm{S}}(z)$ is nondecreasing. Without loss of generality, we may assume that the slope of T^{S} at any point $z \in [0, A] \setminus \operatorname{im} q$ is zero (otherwise replacing the payment at $z \notin \operatorname{im} q$ by the payment for the next highest quantity in $\operatorname{im} q$ also implements the same mechanism because $v' > 0$). On the other hand, for any $z > z'$, we must have $T^{\mathrm{S}}(z) - T^{\mathrm{S}}(z') < c(z - z')$, because $S(z)$ is nondecreasing. Now, suppose that (IC-T) were violated for type θ. In that case, we must have that there exists a $\theta' \neq \theta$ and $z \geq q(\theta')$ with

$$\theta v(z) - t(\theta') - c(z - q(\theta')) > \theta v(q(\theta)) - t(\theta) = \max_{z} \theta v(z) - T^{\mathrm{S}}(z),$$

where the equality follows from the fact that T^{S} implements (q, t). But this implies

$$\theta v(z) - T^{\mathrm{S}}(q(\theta')) - c(z - q(\theta')) > \theta v(z) - T^{\mathrm{S}}(z),$$

or, equivalently,

$$T^{\mathrm{S}}(z) - T^{\mathrm{S}}(q(\theta')) > c(z - q(\theta')),$$

which contradicts our finding above.

To see that *(i)* implies *(iii)*, note that (IC) follows from (IC-T) because the latter implies

$$\theta v(q(\theta)) - t(\theta) \geq V(\theta, \hat{\theta}) \geq \theta v(q(\hat{\theta})) - t(\hat{\theta}).$$

On the other hand, for any $\theta, \theta' \in \Theta$, we have by the concavity of v that

$$\arg\max_{q \geq q(\theta')} \{\theta v(q) - [t(\theta') + c(q - q(\theta'))]\} = \max\{q(\theta'), q^{\mathrm{LF}}(\theta)\},$$

so, if $q(\theta)$ satisfies the (IC-T) constraint, we must have that $q(\theta) \geq q^{\mathrm{LF}}(\theta)$. Thus (IC-T) implies

(LB).

For the converse, suppose that $q(\theta) \geq q^{\mathrm{LF}}(\theta)$. We first establish for the lowest type $\underline{\theta}$ that $t(\underline{\theta}) \leq cq(\underline{\theta})$. To see this, note that (IR) implies

$$t(\underline{\theta}) \leq \underline{\theta} v(q(\underline{\theta})) - \underline{\theta} v(q^{\mathrm{LF}}(\underline{\theta})) + cq^{\mathrm{LF}}(\underline{\theta}),$$

and $\underline{\theta} v(q^{\mathrm{LF}}(\underline{\theta})) - cq^{\mathrm{LF}}(\underline{\theta}) \geq \underline{\theta} v(q(\underline{\theta})) - cq(\underline{\theta})$ by definition of q^{LF}, so $t(\underline{\theta}) \leq cq(\underline{\theta})$. That same argument implies that the *average* price of the units of the good purchased by each type θ is no greater than c (so $t(\theta) \leq cq(\theta)$ for all $\theta \in \Theta$), but to establish that the *marginal* price of any units purchased is no greater than c, note that (IC)—via the Milgrom and Segal (2002) envelope theorem—implies for any two types $\theta, \theta' \in \Theta$ that

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

where the last equality employs integration-by-parts for the Lebesgue-Stieltjes integral. But if $q(\theta) \geq q^{\mathrm{LF}}(\theta)$ for all $\theta \in \Theta$, the concavity of v implies that $v'(q(\theta)) \leq v'(q^{\mathrm{LF}}(\theta)) = c/\theta$, so $t(\theta') - t(\theta) \leq c[q(\theta') - q(\theta)]$. But this implies that the subsidy $S(z)$ is nondecreasing in z for all $z \in \mathrm{im}(q)$, so that, having established *(ii)* $\implies$ *(i)*, we also have *(iii)* $\implies$ *(i)*. □

C.4 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 the only missing detail: 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 $\theta > \hat{\theta}$.

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

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

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

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

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

which implies

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

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

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

Putting these inequalities together, we obtain

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

But then since F is absolutely continuous, we have that $F(\hat{\theta}) - F(\widetilde{\theta}) \sim O(\sqrt{\varepsilon})$ 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.5 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:

$$\max_{(q,t)} \int_{\Theta} \big\{\omega(\theta) \underbrace{[\theta v(q(\theta)) - t(\theta)]}_{\text{consumer surplus}} + \alpha \underbrace{[t(\theta) - cq(\theta)]}_{\text{total profit}} \big\} \, \mathrm{d}F(\theta), \tag{NO-PM}$$

such that (q, t) satisfies (IC), (IR), and (NLS), which can be rewritten as in Appendix A.1 as

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

which differs from (OPT) 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 (IR) constraint in the topping up problem.

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

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

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

$$\int_{\Theta} [\overline{J}(\theta)v(q(\theta)) - cq(\theta)] \ \mathrm{d}F(\theta),$$

so, via similar logic to Appendix A.5, it suffices to show that for any feasible $q \in \mathcal{Q}$,

$$\int_{\Theta} \left[J(\theta) - \overline{J}(\theta)\right][v(q^{\mathrm{SD}}(\theta)) - v(q(\theta))\right] \ \mathrm{d}F(\theta) \geq 0.$$

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

$$\int_{\Theta} \left[J(\theta) - \overline{J}(\theta)\right] v(q^{\mathrm{SD}}(\theta)) \ \mathrm{d}F(\theta) = 0,$$

as we showed in Appendix A.5. Finally, as we also showed in Appendix A.5, because any feasible q is nondecreasing wherever $J \neq \overline{J}$, we have that

$$\int_{\Theta} \left[J(\theta) - \overline{J}(\theta)\right] v(q^{\mathrm{SD}}(\theta)) \ \mathrm{d}F(\theta) \leq 0,$$

completing the proof.

C.6 Comparative Statics Proofs

To prove the comparative statics stated in Section 5.4, 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 5 (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

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

Consider the parametrized objective

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

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

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

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). □

C.6.1 Proof of Proposition 5 (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 $\int_{\theta}^{\overline{\theta}} [\omega(s) - \alpha] \,\mathrm{d}F(s)$ 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 5, we have that $\overline{J|_{\underline{\theta},\theta}}(\theta)$ increases as a function of θ. As a result, the *set* of types with $\overline{J|_{\underline{\theta},\theta}}(\theta) > \theta$ (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 5 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, $S(q)$, 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 11

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

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

C.6.2 Proof of Proposition 6 (Change in Interdependence)

Proof. As discussed in Section 5.4, 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 5 applies, leading to a less generous subsidy mechanism. □

C.6.3 Proof of Proposition 7 (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 $v' : [0, A] \to \mathbb{R}_+$ 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.

□

C.7 Proofs of Extensions

C.7.1 Proof of Proposition 8 (Equilibrium Effects)

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

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

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

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

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

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

subject to $q(\theta) \geq D(p,\theta)$ and $p = S^{-1}\left(\int_\Theta q(\theta) \, \mathrm{d}F(\theta)\right)$. But by the First Welfare Theorem, we have that q^* defined as in Proposition 8 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

$$\int_\Theta [J(\theta)v(q^*(\theta)) - pq^*(\theta)] \, \mathrm{d}F(\theta) + \int_\Theta pq^*(\theta) \, \mathrm{d}F(\theta) - \alpha C\left(\int_\Theta q^*(\theta) \, \mathrm{d}F(\theta)\right)$$
$$\geq$$
$$\int_\Theta [J(\theta)v(q(\theta)) - pq(\theta)] \, \mathrm{d}F(\theta) + \int_\Theta pq(\theta) \, \mathrm{d}F(\theta) - \alpha C\left(\int_\Theta q(\theta) \, \mathrm{d}F(\theta)\right),$$

given our previous finding that

$$\int_\Theta [H(\theta)v(q^*(\theta)) - pq^*(\theta)] \, \mathrm{d}F(\theta) + \int_\Theta pq^*(\theta) \, \mathrm{d}F(\theta) - \alpha C\left(\int_\Theta q^*(\theta) \, \mathrm{d}F(\theta)\right)$$
$$\geq$$
$$\int_\Theta [H(\theta)v(q(\theta)) - pq(\theta)] \, \mathrm{d}F(\theta) + \int_\Theta pq(\theta) \, \mathrm{d}F(\theta) - \alpha C\left(\int_\Theta q(\theta) \, \mathrm{d}F(\theta)\right),$$

which requires only that

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

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

C.7.2 Proof of Proposition 9 (Exogenous Taxation)

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

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 15.

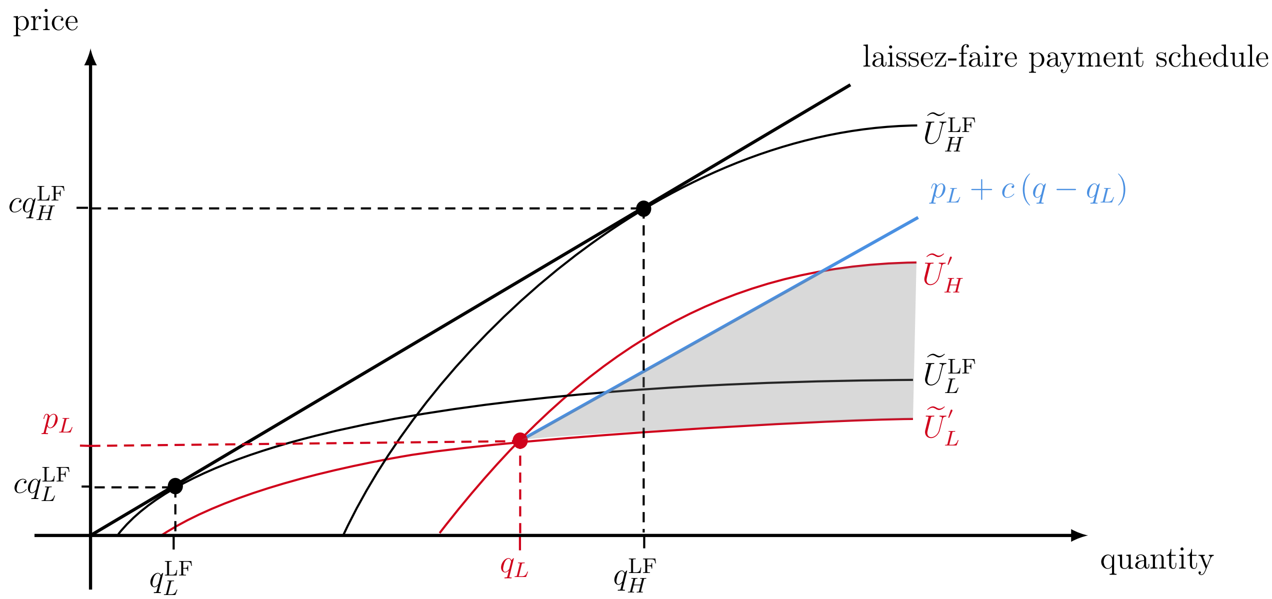

Figure 15: Illustrating the subsidy design problem using indifference curves.

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. As discussed in Section 2.4, the (IR) constraint requires that (p_L, q_L) lies on or below $\widetilde{U}_L^{\text{LF}}$ and that (p_H, q_H) lies on or below $\widetilde{U}_H^{\text{LF}}$, and the (NLS) constraint requires p_L and p_H to be nonnegative. Fixing a choice of (p_L, q_L), the (IC-T) constraint requires (p_H, q_H) to lie in the gray shaded region in Figure 15 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 15: the social planner minimizes the subsidy (the vertical distance between the laissez-faire payment schedule and the price-quantity pair) on $\widetilde{U}_L'$, for example, at q_L^{LF} (as a result of the tangency condition for q_L^{LF} on $\widetilde{U}_L^{\text{LF}}$).

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 tightens the (IC-T) constraint for H-type consumers. 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, $\widetilde{U}_L'$, and the (IC-T) constraints 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.[27] 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 L-type's (IC-T) constraint, reflecting the L-type's incentive to deviate upwards.

[27] 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 tightens the (IC-T) constraint for H-type consumers, 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-T) 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-T) 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 16(a), if $q_L \leq q_H^{\text{LF}}$, the downward (IC-T) 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 16(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 16(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 16(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.

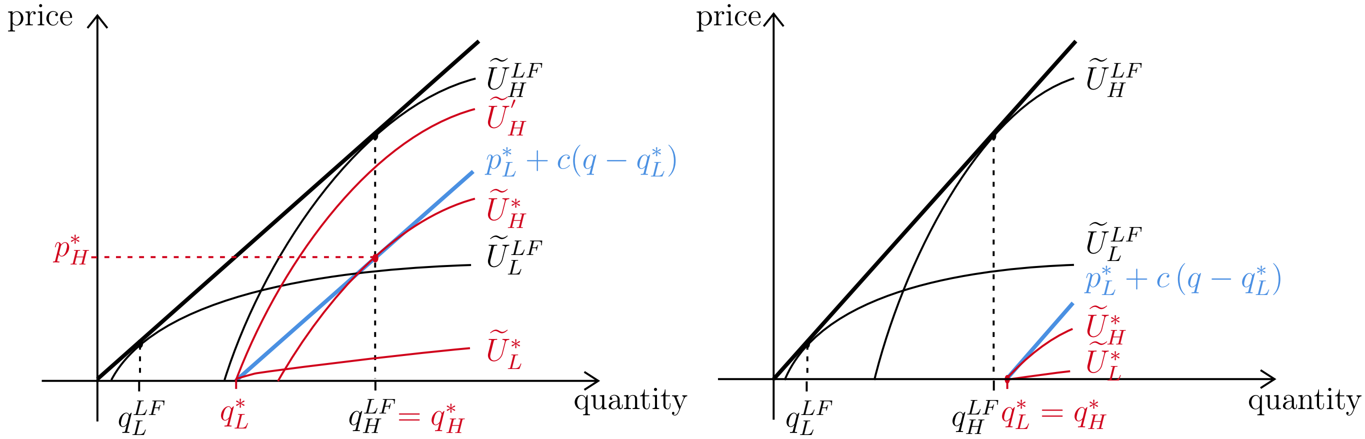

(a) $q_L^* \le q_H^{\text{LF}} = q_H^*$ *(b)* $q_L^* > q_H^*$

Figure 16: 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 $q_L^* > q_L^{\text{LF}}$ 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 $q_L^* = q_L^{\text{LF}}$ and $p_L^* = cq_L^{\text{LF}}$.

The social planner can offer nondistortive cash subsidies to the H-type only until the L-type's upward (IC-T) constraint binds, which happens when $\theta_L v(q_H^{\text{LF}}) - p_H = \theta_L v(q_L^{\text{LF}}) - cq_L$. 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-T) constraint is just binding: $\theta_L v(q_H) - p_H = \theta_L v(q_L^{\text{LF}}) - cq_L$. 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 $p_H^* = \theta_L v(q_H^*) - \theta_L v(q_L^{\text{LF}}) + cq_L$, as illustrated in Figure 17.

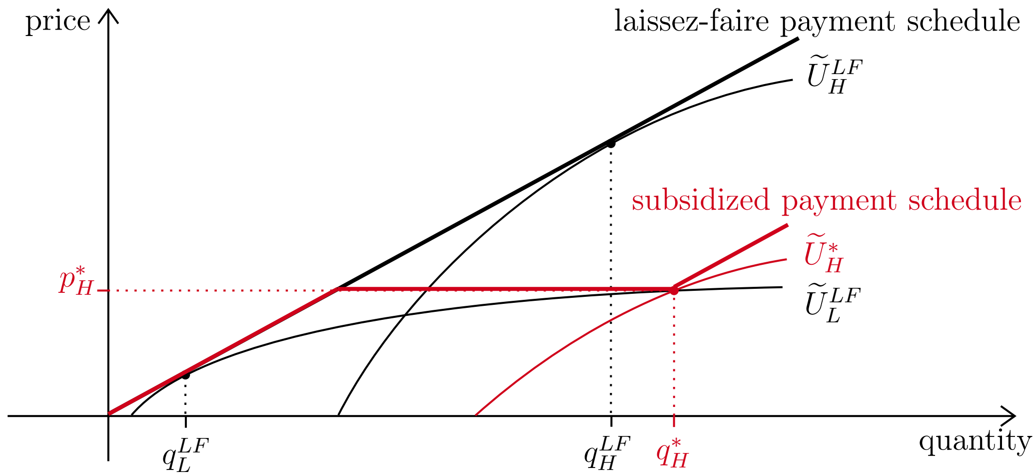

Figure 17: Illustrating the optimal subsidized payment schedule with two types satisfying $\omega_L < \alpha < \omega_H$, with $\mathbf{E}[\omega] \leq \alpha$.

In short, the optimal subsidy mechanism with increasing welfare weights subsidizes consumption of $q_H^* > q_H^{\text{LF}}$ 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-T) 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-T) constraint, the social planner chooses q_L and q_H to maximize the weighted objective

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

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

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

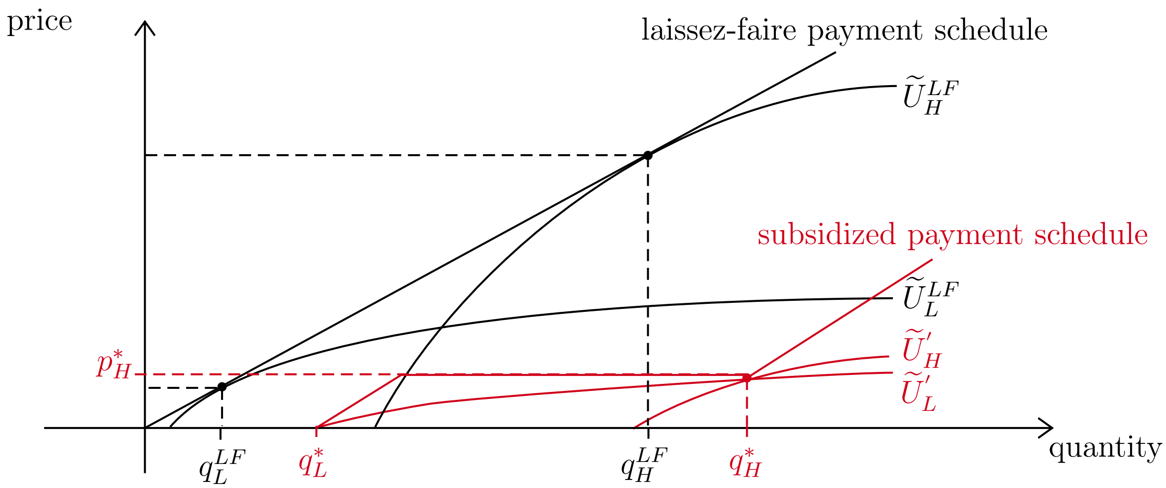

Figure 18: Illustrating an optimal subsidized payment schedule for $\mathbf{E}[\omega] > \alpha$ *and* $\omega_H > \alpha$ *with upward distortion of both types' consumption.*

The optimal mechanism is illustrated in Figure 18.

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-T) 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 toppinng 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 $\int_{\theta}^{\overline{\theta}}[\omega(s)-\alpha]\ \mathrm{d}F(s)$. As illustrated in Figure 19, 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.

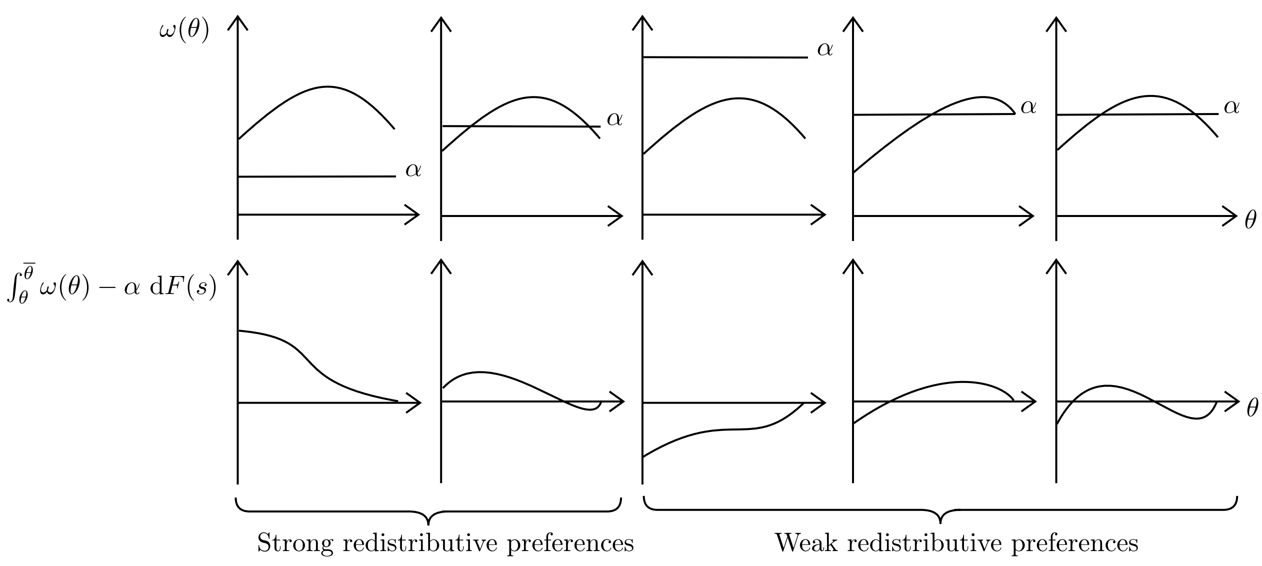

Figure 19: Plots of $\omega(\theta)$ and $\int_{\theta}^{\overline{\theta}}[\omega(s)-\alpha]\ \mathrm{d}F(s)$ for inverted U-shaped welfare weights

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 19, 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 20, 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).

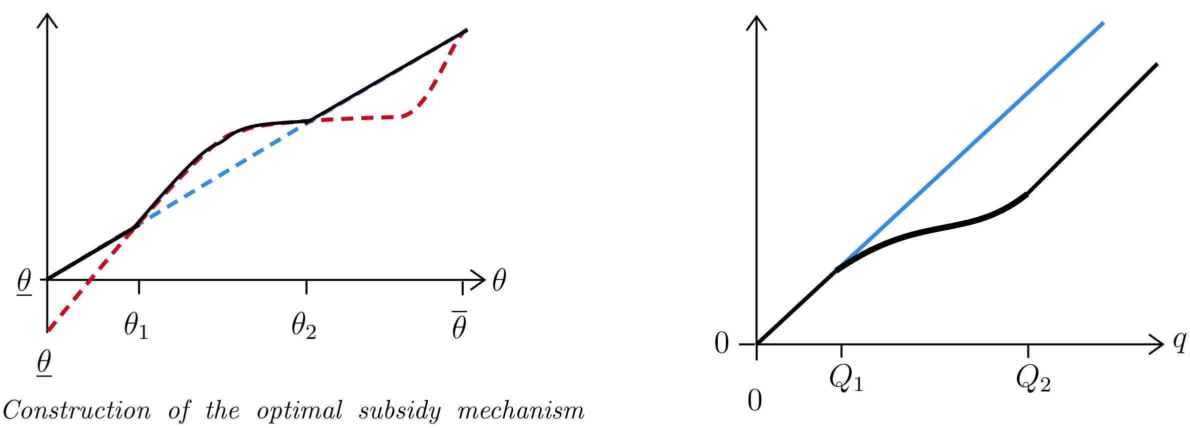

(a) Construction of the optimal subsidy mechanism (virtual welfare in red, lower bound in blue, subsidy type in black)

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

Figure 20: 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 $\int_{\theta}^{\overline{\theta}}[\omega(s)-\alpha]\ \mathrm{d}F(s)$. Once again, as illustrated in Figure 21, 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.

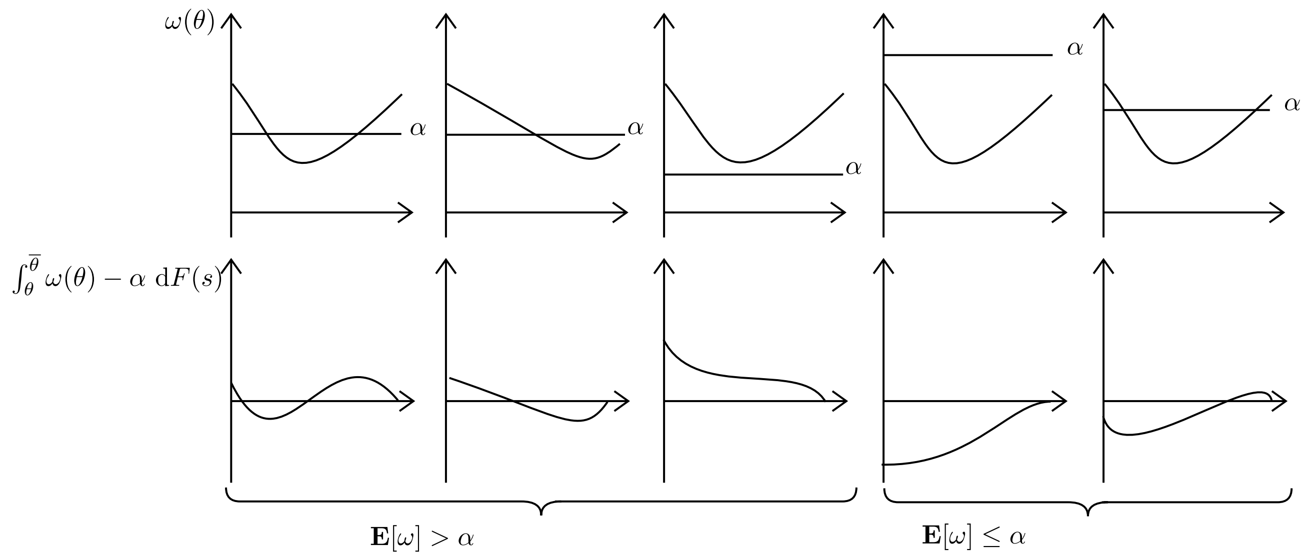

Figure 21: Plots of $\omega(\theta)$ *and* $\int_{\theta}^{\overline{\theta}}[\omega(s)-\alpha]\ \mathrm{d}F(s)$ *for U-shaped welfare weights*

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 22.

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.

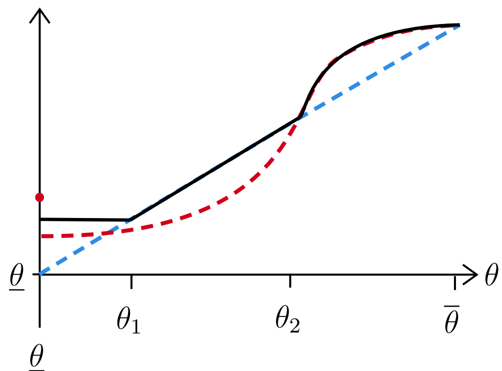

(a) *Construction of the optimal subsidy mechanism (virtual welfare in red, lower bound in blue, subsidy type in black)*

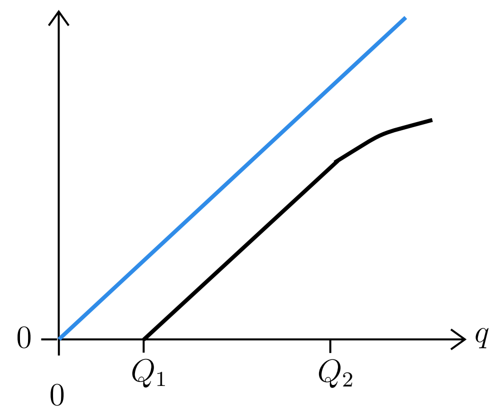

(b) *Payment schedule (laissez-faire in blue, subsidized in black)*

Figure 22: 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 $v : [0,1] \to [0,1]$ be defined by $v(q) = q$, where q is interpreted as the *probability* that the consumer is allocated the good.[28] We suppose, moreover, that $c \in \operatorname{int} \Theta$, so not all agents purchase the good in the laissez-faire economy.

The optimal subsidy mechanism satisfies the *linear* program

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

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

[28] 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 12

Proposition 12 (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*

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

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

$$c = p + \frac{\int_p^{\overline{\theta}} \omega(s) - \alpha \, \mathrm{d}F(s)}{\alpha f(p)}.$$

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

(i) $\int_c^{\theta} [J(s) - c] \, \mathrm{d}F(s) \leq 0$ *for all* $\theta \geq c$, *and*

(ii) $\int_{\theta}^{c} [J(s) - c] \, \mathrm{d}F(s) \leq 0$ *for all* $\theta \leq c$, *with equality for some* $\theta < c$.

In that case, writing $\theta_1 < \theta_2 < ...\theta_I := c$ for the set of types less than c for which $\int_{\theta}^{c} [J(s) - c] \, \mathrm{d}F(s) = 0$, 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 $i = 1, 2, ..., I-1$, 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 12 follows as a direct application of the solution of general linear programs with lower-bound constraints, presented in Proposition 13 below. The two outstanding details concern the price of the subsidized good and are as follows:

(a) The claim that $c = p + \frac{\int_p^{\overline{\theta}} \omega(s) - \alpha \, \mathrm{d}F(s)}{\alpha f(p)}$ for any optimal subsidy price $p < c$. This is the first-order condition for the maximization problem $\max_{p \leq c} \int_p^c [J(\theta) - c] \, \mathrm{d}F(\theta)$, 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 $\int_p^c [J(\theta) - c]\ \mathrm{d}F(\theta)$ 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

$$\max_{q:\Theta\to[0,A]} \int_\Theta \phi(\theta) q(\theta)\ \mathrm{d}F(\theta),\ \text{subject to } \forall\theta\in\Theta:\ q(\theta) \geq q^{\mathrm{L}}(\theta) \qquad \text{(LP)}$$

for a monotone function $q^{\mathrm{L}} : \Theta \to [0, A]$ and an integrable $\phi : \Theta \to \mathbb{R}$, 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 $\overline{\theta}$, 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 $\Phi(\theta) = \int_{\underline{\theta}}^{\theta} \phi(s)\ \mathrm{d}F(s)$ and let $\operatorname{decmin}\Phi$ be the *decreasing minorant* of Φ, which is the largest nonincreasing function everywhere below Φ, as illustrated in Figure 23. 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) $\Phi(\theta) = \operatorname{decmin}\Phi(\theta)$ for all $\theta \in \Theta^i$, or (b) $\Phi(\theta) > \operatorname{decmin}\Phi(\theta)$ for all $\theta \in \operatorname{int}\Theta^i$ (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 23.

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

Proposition 13 (solutions to linear program)**.** *Any solution of* (LP) *satisfies the following:*

(i) If $\Phi(\theta) = \operatorname{decmin}\Phi(\theta)$ *on* Θ^i*, then* $q^*(\theta) = q^{\mathrm{L}}(\theta)$ *for all* $\theta \in \Theta^i$*.*

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

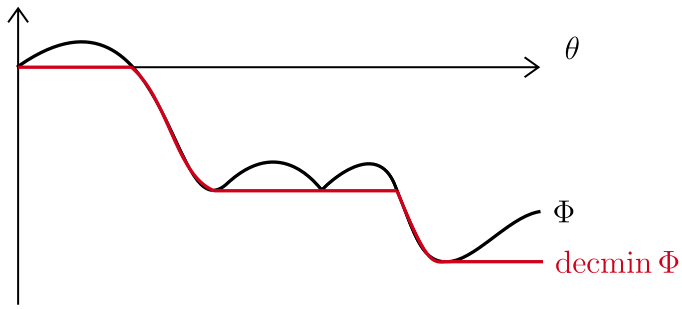

Figure 23: Constructing the Decreasing Minorant

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

While the statement of Proposition 13 appears complicated, the intuitive description of the optimizers is simpler: wherever $\Phi(\theta) = \operatorname{decmin}\Phi(\theta)$, 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 $\Phi > \operatorname{decmin}\Phi$, 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):

$$\begin{aligned}
\int_\Theta \phi(\theta)q(\theta)\ \mathrm{d}F(\theta) &= q(\overline{\theta})\Phi(\overline{\theta}) - \int_\Theta \Phi(\theta)\ \mathrm{d}q(\theta) \\
&\leq \Phi(\overline{\theta}) - \int_\Theta \operatorname{decmin}\Phi(\theta)\ \mathrm{d}q(\theta) \\
&\leq \Phi(\overline{\theta}) - \int_\Theta \operatorname{decmin}\Phi(\theta)\ \mathrm{d}q^{\mathrm{L}}(\theta) \\
&= \Phi(\overline{\theta}) - \operatorname{decmin}\Phi(\overline{\theta}) + \int_\Theta \frac{\mathrm{d}}{\mathrm{d}\theta}\operatorname{decmin}\Phi(\theta)q^{\mathrm{L}}(\theta)\ \mathrm{d}\theta.
\end{aligned}$$

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

We now argue that the upper bound is obtained by any solution taking the form described in Proposition 13 and that any other feasible q results in a strictly lower objective value. We calculate $\int_\Theta \phi(\theta)q^*(\theta)\ \mathrm{d}F(\theta)$ separately for each interval of type (i), (ii), and (iii) described in Proposition 13.

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

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

because $q^*(\theta) = q^{\mathrm{L}}(\theta)$ on Θ^i and because $\operatorname{decmin}\Theta(\theta) = \Theta(\theta)$. 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

$$\int_{\Theta^i} \phi(\theta)q^*(\theta)\ \mathrm{d}F(\theta) = \int_{\Theta^i} \phi(\theta)\ \mathrm{d}F(\theta) = \Phi(\overline{\theta}) - \operatorname{decmin}\Phi(\overline{\theta}) + \int_{\Theta^i} \frac{\mathrm{d}}{\mathrm{d}\theta} \operatorname{decmin}\Phi(\theta)q^{\mathrm{L}}(\theta)\ \mathrm{d}\theta,$$

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

$$\int_{\Theta_i} \phi(\theta)q(\theta)\ \mathrm{d}F(\theta) = \Phi(\overline{\theta}) - \Phi(\theta_-) - \int_{\Theta_i} (\Phi(\theta) - \Phi(\theta_-))\ \mathrm{d}q(\theta),$$

where θ_- is the left endpoint of Θ^i so $\operatorname{decmin}\Phi(\overline{\theta}) = \Phi(\theta_-)$ 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

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

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

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

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

□