The Public Option and Optimal Redistribution*

Zi Yang Kang†

This version: January 2023
Latest version: available here

Abstract

This paper examines how the equilibrium effects of a public option on the private market impact its optimal design. I develop a model in which a policymaker can choose the quality and allocation of the public option, which affect the prices of private goods (and vice versa) in equilibrium. I demonstrate how these equilibrium effects change both the optimal quality and optimal allocation: they create new incentives to distort quality in either direction depending on the policymaker's redistributive objective and provide a new justification for rationing the public option rather than using market-clearing prices. Finally, I show how my results can accommodate additional frictions in the private market and additional policy instruments.

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

Keywords: public option, redistribution, inequality, mechanism design

---

* I am especially grateful to Paul Milgrom, Andy Skrzypacz, and Shosh Vasserman for many helpful discussions. I also thank Mohammad Akbarpour, Morris Ang, Lanier Benkard, Jeremy Bulow, Cody Cook, José Ignacio Cuesta, Rebecca Diamond, Piotr Dworczak, Liran Einav, Matt Gentzkow, Ben Hébert, Ravi Jagadeesan, Diego Jiménez-Hernández, Ellen Muir, Mike Ostrovsky, Peter Reiss, Claudia Robles-Garcia, Ilya Segal, Paulo Somaini, Juan Carlos Suárez Serrato, Winnie van Dijk, Mitchell Watt, Frank Yang, and Ali Yurukoglu for comments and suggestions. The abstract of an earlier version of this paper appeared in the *Proceedings of the 22nd ACM Conference on Economics and Computation (EC'21)* under the title "Optimal Public Provision of Private Goods." This paper has also benefitted from numerous conference and seminar participants at EC'21 and Stanford. Finally, I am grateful to the SIEPR Program in Regulatory Policy for financial support.
† Graduate School of Business, Stanford University; zykang@stanford.edu.

1 Introduction

Governments frequently redistribute by supplying a public alternative, or a "public option," that competes with goods sold by private producers. Public housing programs, for example, allow eligible individuals to rent affordable housing units at lower prices than private apartments of similar quality. Likewise, individuals may choose to attend either a public or private school and to seek either public or private health care.

While it is easy to see how the public option directly benefits those who consume it, the public option might also indirectly affect those who do not, due to its equilibrium impact on the prices of private goods. On one hand, advocates argue that the public option could exert downward pressure on the prices of private goods. For example, in their book on *The Public Option*, Sitaraman and Alstott (2019) write:

> "[T]he presence of the public option puts pressure on private actors to provide better service at lower cost."

On the other hand, critics contest that the public option might exert upward pressure on the prices of private goods. For instance, in his book on *Capitalism and Freedom*, Friedman (1962) cautions:

> "Far from improving the housing of the poor, as its proponents expected, public housing has done just the reverse...Some families have probably been better housed than they would otherwise have been—those who were fortunate enough to get occupancy of the publicly built units. But this has only made the problem for the rest all the worse, since the average density of all together went up."

Yet, despite the importance of these equilibrium effects to policy debates about the public option, standard models in mechanism design and theoretical public finance do not allow for them.

In this paper, I study how the equilibrium effects of a public option on private goods change its optimal design, namely, how it should be allocated and at what quality it should be provided. I show that these equilibrium effects qualitatively change the nature of the optimal allocation: they provide a new justification for rationing the public option—for example, through the use of lotteries. These equilibrium effects also generate a novel distortion in the optimal quality of the public option, the sign of which depends on the policymaker's redistribution objective.

To study these equilibrium effects, I develop a model of a policymaker who can provide a public option with accompanying cash transfers in a given market. The market consists of risk-neutral consumers, each of whom demands a single unit of an indivisible good, and private producers who supply it competitively at different quality levels: for instance, apartments might differ in size and schools, in teacher–student ratio. Consumers are distinguished by two characteristics. First, consumers have heterogeneous consumption preferences over quality, which determine their marginal utility for quality. Second, consumers have heterogeneous welfare weights in the social welfare function, which captures the policymaker's redistribution objective. A consumer's welfare weight measures the social value of giving him one unit of money. For example, a consumer might have a high welfare weight if he is poor, infirm, or belongs to a socially disadvantaged group. As a common requirement of many public options (e.g., public education and health care) is equal access, the policymaker cannot screen consumers on the basis of their welfare weights. Even when equal access is not assumed, determinants of a consumer's welfare weight, such as his health or expected future income, might be private information. However, the correlation between consumption preferences and welfare weights enables the policymaker to infer—albeit imperfectly—a consumer's welfare weight from his consumption behavior.

The objective of the policymaker is to maximize the social welfare function, namely, expected total weighted surplus. To this end, she chooses a quality level at which to supply the public option and an incentive-compatible and individually rational allocation mechanism to allocate the public option. This formulation allows the policymaker to randomly allocate, or "ration," the public option. In particular, it accommodates lotteries and waiting lists, which are common in the allocation of public housing and health care. Consumers can—but are not required to—apply for the public option. Consumers who are not allocated the public option proceed to the private market, where the competitive equilibrium is realized.

A key feature of my model is that the allocation of the public option can affect the prices of private goods. This is motivated by empirical evidence, which shows how this effect can be large in some markets (Currie and Gahvari, 2008). To capture this, goods of higher quality in my model require a larger amount of an input in their production. In turn, input has an upward-sloping supply curve. For example, larger apartments require more space and higher teacher–student ratios require more teachers, where space and teachers are in scarce supply. One one hand, the public option reduces the residual supply of input for private goods: public housing takes up space that could otherwise have been used for private apartments, just as public schools employ teachers who could otherwise have taught in a private school. On the other hand, the public option also reduces the residual demand for private goods: consumers who receive public housing no longer rent private apartments, just as students in public schools no longer consume private education. In equilibrium, the public option can therefore either raise or lower the prices of private goods.

To analyze how these equilibrium effects affect the optimal design of the public option, I decompose the welfare impact of a public option into two components, each of which captures a fundamental economic force in the policymaker's problem. While the *direct effect* of the public option is equal to the change in social welfare holding fixed the prices of private goods, the *indirect effect* of the public option is equal to the change in social welfare arising from changes in the prices of private goods. From an economic perspective, while the direct effect measures how well the policymaker can screen consumers, the indirect effect measures how much pecuniary externality the public option exerts on the private market.

My first main result shows that the trade-off between the direct and indirect effects of the public option qualitatively changes the optimal allocation: it provides a new justification for rationing. To understand the underlying intuition, consider an example in which the policymaker can supply a thousand units of public housing at a given quality without affecting prices of apartments, but she significantly raises the price of input if she supplies more than a thousand units. This large and negative indirect effect prevents her from supplying more than a thousand units of public housing. Conditional on supplying a thousand units, a market-clearing price selectively allocates public housing units to richer consumers, while rationing offers poorer consumers the opportunity to derive some direct benefit from public housing as well. Therefore, rationing is optimal in this example if the policymaker places sufficiently high welfare weight on the poorest consumers.

Conversely, the optimal allocation does not require rationing when there is no trade-off between the direct and indirect effects of the public option. This is the case when input supply is perfectly elastic (e.g., there is abundant space for housing), so that allocation of the public option has no impact on prices in the private market. To understand why, suppose that the policymaker rations some consumers. Then there must be some (non-negative) gain in the direct effect from allocating to these consumers with slightly higher probability: if not, then the policymaker would not have allocated the public option to these consumers with positive probability to begin with. Because the supply of input is perfectly elastic, the new policy entails no loss in the indirect effect and does not affect incentive constraints: the allocation of the public option does not change the prices of private goods. Thus, extending this logic, the policymaker could have allocated the public option with probability one to these consumers in the first place, which removes any need for rationing.

My second main result shows that the trade-off between the direct and indirect effects of the public option also changes the optimal quality: it provides a new incentive for the policymaker to either raise or lower the quality of the public option. Each effect introduces a quality distortion. On one hand, the direct effect introduces a screening distortion, which arises because consumers have private information about their consumption types and the policymaker has a redistributive (non-utilitarian) objective. The screening distortion is always negative, which captures the well-known intuition that a low-quality public option allows the policymaker to better target poorer consumers. On the other hand, the indirect effect introduces a pecuniary externality distortion, which is new. Intuitively, this distortion arises because increasing the quality of the public option results in a first-order increase in the prices of private goods; in turn, this has a positive indirect effect on private producers but a negative indirect effect on consumers and the policymaker. Thus, depending on the relative weights of these agents in the policymaker's objective function, she faces additional incentives to either raise or lower the quality of the public option.

My results generalize in three different dimensions that have featured in policy debates about the public option. First, I show how externalities and paternalistic preferences—which result in different weights on utility and consumption for each consumer in the social welfare function—can be incorporated into my framework. Second, I show how market power in the private market can be accommodated. Third, I show how my analysis extends when the policymaker has access to a richer set of policy instruments than a public option, such as taxes and subsidies on quality.

1.1 Related literature

The main feature of this paper that distinguishes it from much of the existing literature on in-kind transfers is the ability of the public option to affect the prices of private goods. Since the pioneering work of Nichols and Zeckhauser (1982), there has been an extensive literature that studies the efficiency and redistributive impact of in-kind transfers. This literature has shown how in-kind transfers can screen consumers better than cash transfers (e.g., Blackorby and Donaldson, 1988) and how a private market might impact screening (e.g., Besley and Coate, 1991; Gahvari and Mattos, 2007). However, the public option does not affect the prices of private goods these models as input supply is assumed to be perfectly elastic. Consequently, the results of these papers speak to only the direct effect of the public option in my framework. By contrast, I develop a model that allows the public option to affect the prices of private goods, which enables me to study the trade-off between the direct and indirect effects.

In addition, the ability of the public option to affect the prices of private goods also distinguishes this paper from the growing literature on redistributive mechanism design. This literature stems from the seminal work of Weitzman (1977), who observed that rationing can help redistribute a good to consumers when the policymaker seeks to maximize a different objective than utilitarian welfare. This literature has formalized Weitzman's observation using the tools of mechanism design and characterized optimal mechanisms in general settings (e.g., Condorelli, 2013; Che, Gale, and

Kim, 2013b; Dworczak ⓡ Kominers ⓡ Akbarpour, 2021; Akbarpour ⓡ Dworczak ⓡ Kominers, 2022). 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, 2022; Akbarpour ⓡ Budish ⓡ Dworczak ⓡ Kominers, 2021; Pai and Strack, 2022). However, these models normalize consumers' outside options to zero by focusing on settings without a private market. By contrast, I show that this normalization entails a loss of generality. In particular, I show that the ability of the public option to affect the prices of private goods affects the policymaker's optimal choices of quality and allocation for the public option.

The presence of a private market in my model connects this 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), and optimal contracting between firms (e.g., Calzolari and Denicolò, 2015; Kang and Muir, 2022). My paper contributes to this literature by enriching the private market in three ways. First, consumption in my model generates pecuniary externalities on other consumers in the private market. Second, consumers have heterogeneous welfare weights that affect the policymaker's preferences over private market consumption. Third, I show how my results extend to the case of an imperfectly competitive market. These features of the private market are important to capture in my setting of redistribution via the public option, but also complicate the interaction between the public option and private market.

A growing number of empirical papers estimate the effect of public programs on private good prices, therefore motivating and justifying the premise of my analysis. For example, in the context of housing, Diamond and McQuade (2019) find that housing financed by the Low Income Housing Tax Credit (LIHTC) increased house prices by 6.5% in low-income neighborhoods but decreased house prices by 2.5% in high-income neighborhoods. Baum-Snow and Marion (2009) also find that LIHTC-financed housing increases house prices in low-income neighborhoods. In the context of health care, Atal, Cuesta, González, and Otero (2021) find that public pharmacies in Chile induce a 1.1% increase in private pharmacy prices. In the context of education, Dinerstein and Smith (2021) show that the supply of private schools in New York City is considerably elastic: local private school supply fell by 6% for every $1,000 in projected funding increase per student in a public school. In the context of food and nutrition, Handbury and Moshary (2021) show that grocery retail chains lower prices in response to the expansion of school lunch programs. These findings suggest that the ability of the public option to affect the prices of private goods is an important modeling feature to capture.

This paper also complements a growing empirical literature that studies the provision of public options in different markets. These markets include housing (e.g., Waldinger, 2021; van Dijk, 2019), education (e.g., Epple and Romano, 1998; Hoxby, 2000; Dinerstein and Smith, 2021; Dinerstein, Neilson, and Otero, 2020), health care (e.g., Duggan and Scott Morton, 2006; Curto, Einav, Finkelstein, Levin, and Bhattacharya, 2019; Atal et al., 2021), and food and nutrition (e.g., Jiménez-Hernández and Seira, 2021). My paper complements this literature by providing a general theoretical framework for determining how and when to redistribute via a public option. In turn, this sheds light on what markets are more favorable to redistribution via a public option.

My results on the optimality of rationing relate my paper to an extensive mechanism design literature. In this literature, rationing (or random allocation) usually arises for two reasons. First, non-concave capacity costs, often modeled as a capacity constraint, may lead to rationing in optimal pricing (e.g., Hotelling, 1931; Myerson, 1981; Maskin and Riley, 1984; Wilson, 1988; Bulow and Roberts, 1989; Loertscher and Muir, 2022) and redistribution (e.g., Weitzman, 1977; Condorelli, 2013; Dworczak ⓡ al., 2021; Akbarpour ⓡ al., 2021). Second, budget constraints may lead to rationing in optimal intervention due to adverse selection (e.g., Samuelson, 1984) and redistribution (e.g., Dworczak ⓡ al., 2021). However, the policymaker in my model faces neither capacity nor budget constraints. Rather, rationing arises for a novel reason—namely, the trade-off between the direct and indirect effects of the public option.

From a methodological perspective, my results on rationing build on techniques developed in mechanism design and information design. In particular, "generalized ironing" and concavification techniques à la Myerson (1981), Aumann and Maschler (1995), and Kamenica and Gentzkow (2011) have been widely used to characterize optimal mechanisms (e.g., Toikka, 2011; Hartline, 2013; Dworczak ⓡ al., 2021; Loertscher and Muir, 2022). However, these techniques do not generally apply in my setting because the policymaker does not face a standard type of constraint (i.e., either a capacity or budget constraint). Consequently, my characterization of the optimal mechanism is non-constructive and exploits the mathematical results of Bauer (1958) and Szapiel (1975). This approach is closer to the extreme-point approach taken in the literature (e.g., Skreta, 2006; Manelli and Vincent, 2007; Kleiner, Moldovanu, and Strack, 2021; Nikzad, 2022). As the form of constraint that arises endogenously in my model are affine in the choice variable (i.e., the allocation function), it can be characterized with a Carathéodory-like theorem (e.g., Le Treust and Tomala, 2019; Doval and Skreta, 2022). Despite this non-constructive characterization of the optimal mechanism, I develop a method to explicitly compute the optimal mechanism using tools from infinite-dimensional quadratic programming (e.g., Reid, 1968; Barron, 1983), which has recently been used to characterize optimal mechanisms in a different setting (Kang, 2022).

Finally, this paper shares a number of ideas in common with a companion paper, Kang (2022), although the modeling assumptions and results are different. In that paper, I study the optimal regulation of an externality that cannot be taxed (e.g., due to the inability to measure it directly). Instead, the policymaker indirectly taxes a proxy good, the consumption of which is correlated with how much externality each consumer produces. I show that nonlinear taxation is optimal and solve for the optimal mechanism. In this paper, the consumption of the good can be viewed as a proxy for each consumer's welfare weight. Although there is no real externality in the baseline version of my model, there is a pecuniary externality that each consumer's consumption exerts on others in the private market. When welfare weights are heterogeneous, this pecuniary externality, like a real externality, warrants intervention. The two papers differ in the set of policy instruments available to the policymaker, which leads to different results. The policymaker faces a restricted set of policy instruments in this paper: she can neither control the private market nor set a nonlinear tax on quality. This restriction is motivated by practical constraints in markets that prevent the policymaker from perfectly regulating the private market or accurately measuring quality. This restriction leads to different results obtained by the two papers: rather than nonlinear taxation, I show that the policymaker rations the public option as a second-best solution. Also, she trades off direct effects of the public option with indirect effects that arise in equilibrium; this is a novel trade-off that arises only because the policymaker does not control the entire market.

1.2 Organization of this paper

The remainder of this paper is organized as follows. Section 2 develops a model of public provision, while Section 3 shows how the policymaker's problem can be analyzed as a constrained mechanism design problem. Section 4 characterizes optimal mechanisms, and Section 5 studies the optimal quality. Section 6 then examines extensions, and Section 7 concludes. Omitted proofs are provided in Appendix A; and Appendix B and Appendix C contain additional results and discussion.

2 Model

In this section, I develop a model of redistribution in which the public option can affect the prices of private goods. I then formulate the mechanism design problem that the policymaker faces.

2.1 Setup

There is a unit mass of risk-neutral consumers in a private market for an indivisible good, which is available at different levels of quality $q \in \mathbb{R}_+$. Consumers have quasilinear utility over money and demand only a single unit of the good. Each consumer is distinguished by his consumption type θ, which determines the utility $u(q, \theta)$ that he derives from consuming a unit of good that has quality q. The distribution of θ is denoted by F and is assumed to have positive density f on the compact domain $[\underline{\theta}, \overline{\theta}] \subseteq \mathbb{R}_+$. I impose the following assumption on consumers' utility functions:

Assumption 1. *The utility function $u : \mathbb{R}_+ \times [\underline{\theta}, \overline{\theta}] \to \mathbb{R}$ is twice continuously differentiable and satisfies, for $q, \theta > 0$:* (i) *$\partial u/\partial q > 0$;* (ii) *$\partial^2 u/\partial q^2 < 0$;* (iii) *$\partial^2 u/\partial q \, \partial \theta > 0$; and* (iv) *$\partial u/\partial q \to +\infty$ as $q \to 0$ and $\partial u/\partial q \to 0$ as $q \to +\infty$.*

The conditions in Assumption 1 are weak and standard. In particular, conditions (i) and (ii) require that consumers derive positive but diminishing marginal utility from quality consumption. The utility function satisfies the strict single-crossing property—condition (iii)—so that consumers with higher consumption types derive higher marginal utility from quality consumption. Finally, condition (iv) is an Inada condition, which requires that consumers derive high marginal utility on initial units of quality and low marginal utility when quality is sufficiently abundant.

The good is supplied competitively by producers in a private market, where production of the good at quality q requires q units of a scarce input. For example, the quality of an apartment might depend on its size, which requires space as an input; likewise, the quality of a school might depend on teacher–student ratio, which requires teachers as an input. The input market is competitive.

The cost of supplying each quality level of the private good thus depends on the price of each unit of input. This allows me to capture the pecuniary externality that each consumer's quality consumption exerts on others. For example, by consuming larger apartments, consumers drive up the marginal price of land at the expense of others; similarly, consuming a higher teacher–student ratio makes it more expensive to hire teachers.

The cost of supplying each unit of good at quality q is denoted by $c(q) + q \cdot p(Q)$. Here, c denotes the cost function for converting q units of input into the good, p denotes the inverse supply function of the input, and Q denotes the total amount of input required. Notice that Q is equal to the aggregate amount of quality consumed in the market; consequently, if each consumer with consumption type θ consumes a good of quality $q(\theta)$, aggregate quality is given by

$$Q = \int_{\underline{\theta}}^{\overline{\theta}} q(\theta) \, \mathrm{d}F(\theta).$$

I impose the following assumption on the cost function and the inverse supply function:

Assumption 2. *The cost function* $c : \mathbb{R}_+ \to \mathbb{R}_+$ *is twice continuously differentiable and satisfies:* (i) $\partial c/\partial q \geq 0$ *and* (ii) $\partial^2 c/\partial q^2 \geq 0$*. The inverse supply function* $p : \mathbb{R}_+ \to \mathbb{R}_{++}$ *is continuously differentiable and satisfies* (iii) $\partial p/\partial Q \geq 0$.

The conditions in Assumption 2 are also weak and standard. The cost of transforming input into each good is non-decreasing and convex in its quality, as required by conditions (i) and (ii). Moreover, condition (iii) stipulates that the supply function for input is weakly upward sloping, so that the inverse supply function p is non-decreasing.

Assumption 2 thus generalizes the assumption of linear conversion between input and quality commonly imposed in the literature. Besley and Coate (1991) and Gahvari and Mattos (2007), for example, assume that producers have a technology that converts pq units of money into a good of quality q. As p is independent of aggregate quality Q, input supply is perfectly elastic; conversion of input into quality is also costless, so $c \equiv 0$. Piazzesi and Schneider (2016) document the use of similar assumptions in the macroeconomics literature on housing. By imposing Assumption 2, I consider the more general case where input supply might have finite elasticity, and where conversion of input into quality might be costly.

Finally, while I have both assumed that quality is unidimensional and abstracted away from non-pecuniary externalities, these can be incorporated into the model. In Section 6, I show how my analysis extends to a model with externalities. Such a model is also mathematically equivalent to one in which quality is multidimensional.

2.2 Laissez-faire equilibrium

I now determine the laissez-faire equilibrium in the private market. As the market is competitive, the price of a good that has quality q is simply equal to its cost, $c(q) + q \cdot p(Q)$. For any level of aggregate quality Q, a consumer's indirect utility function $v_0(\theta, Q)$ solves his utility maximization problem:

$$v_0(\theta, Q) := \max_{q \in \mathbb{R}_+} \left[u(q, \theta) - c(q) - q \cdot p(Q) \right].$$

Let $D(\cdot, \theta)$ denote the inverse of the marginal utility for a consumer with consumption type θ:

$$D(p, \theta) := \left(\frac{\partial u}{\partial q} \right)^{-1} (p, \theta).$$

Thus $D(p,\theta)$ denotes the quality demanded by a consumer with consumption type θ at a marginal price of p per unit of quality. Assumptions 1 and 2 then guarantee that the solution $q_0(\theta,Q)$ for any consumer with $\theta > 0$ is uniquely pinned down by the first-order condition

$$q_0(\theta,Q) = D\left(c'(q_0(\theta,Q)) + p(Q),\theta\right).$$

The equilibrium level of aggregate utility Q_0 is determined by averaging $q_0(\theta,Q)$ over consumption types θ. The following proposition summarizes this discussion.

Proposition 1. *There exists a unique laissez-faire equilibrium under which any consumer with consumption type $\theta > 0$ consumes a good of quality $q_0(\theta,Q_0)$, such that*

$$q_0(\theta,Q_0) = D\left(c'(q_0(\theta,Q_0)) + p(Q_0),\theta\right), \quad \textit{where } Q_0 = \int_{\underline{\theta}}^{\overline{\theta}} q_0(\theta,Q_0)\ \mathrm{d}F(\theta).$$

2.3 Policy design

To complete the model, I formulate the design problem faced by a policymaker who supplies a public option. The policymaker chooses a *policy* (δ, X, T), which comprises:

(i) a quality level $\delta \in \mathbb{R}_+$ at which to supply the public option;

(ii) a mechanism (X,T), consisting of an allocation function $X : [\underline{\theta},\overline{\theta}] \to [0,1]$ and a payment function $T : [\underline{\theta},\overline{\theta}] \to \mathbb{R}$.

Here, the policymaker uses a direct mechanism to allocate the public option, which is without loss of generality by the revelation principle. Under the mechanism (X,T), each consumer truthfully reports his consumption type θ, receives the public option with probability $X(\theta)$, and makes an expected payment of $T(\theta)$ to the policymaker. By setting $(X,T) \equiv (0,0)$, the policymaker can always choose not to supply the public option.

This formulation assumes that the policymaker cannot randomize between different policies. As I show in Section 3, this assumption entails no loss of generality for the purpose of solving for the optimal policy.

This formulation also assumes that the policymaker can choose only a single level of quality for the public option. This assumption is made for simplicity: a single level of quality captures the idea that the public option is "standardized" in order to provide a baseline version of the good. This is not only a good approximation in many existing settings, but also a prevalent view in policy debates of the role that many potential public options might serve (Sitaraman and Alstott, 2019). This assumption can be further justified, for example, by high fixed costs of supplying the public option at additional quality levels.

The policymaker faces a weakly higher cost of supplying the public option than a private producer for a good of the same quality δ. Specifically, the policymaker's cost of supplying each unit of the public option at a quality level δ is equal to $c(\delta) + \kappa + \delta \cdot p(Q)$, where $\kappa \geq 0$. The difference κ captures the policymaker's relative inefficiency at supplying the good as well as any potential costs of administering the public program.

I now describe the timing of the game. First, the policymaker chooses a policy (δ, X, T), after which consumers voluntarily apply for the public option. Second, allocations for the public option are realized, and allocated consumers leave the market. Third, the remaining consumers proceed to the private market, where the competitive equilibrium is realized.

The key feature of this model is the interaction between the mechanism and the private market. There are two ways in which they interact. First, the policymaker's choice of allocation function X affects the aggregate quality of all goods consumed; in turn, this changes the residual supply of input available for private goods. Second, the policymaker's choices of quality δ and allocation function X affect not only how many, but also which, consumers proceed to the private market, and hence the residual demand for private goods. Together, these imply that private market outcomes depends on the policymaker's choices of δ and X. To explicitly capture this dependence, I denote the equilibrium aggregate quality by $Q(\delta, X)$.

Next, I describe the incentive constraints that the policymaker faces. Because the market is large, individual misreports of consumption types do not affect the equilibrium aggregate quality $Q(\delta, X)$. Thus the incentive compatibility constraint, which ensures that each consumer reports his consumption type truthfully, is

$$\theta \in \operatorname*{arg\,max}_{\theta' \in [\underline{\theta}, \overline{\theta}]} \{u(\delta, \theta) \cdot X(\theta') + v_0(\theta, Q(\delta, X)) \cdot [1 - X(\theta')] - T(\theta')\} \quad \text{for all } \theta \in [\underline{\theta}, \overline{\theta}]. \tag{IC}$$

Similarly, the individual rationality constraint, which ensures that each consumer can do no worse by participating in the mechanism, is

$$u(\delta, \theta) \cdot X(\theta) + v_0(\theta, Q(\delta, X)) \cdot [1 - X(\theta)] - T(\theta) \geq v_0(\theta, Q(\delta, X)) \quad \text{for all } \theta \in [\underline{\theta}, \overline{\theta}]. \tag{IR}$$

The incentive constraints (IC) and (IR) reveal the key challenge in the policymaker's problem, namely the nonlinear dependence of consumers' expected utilities on the allocation function X

via aggregate quality $Q(\delta, X)$. Models in the mechanism design literature typically assume that each consumer's expected utility function satisfies the strict single-crossing property in (θ, X), in which case standard arguments imply that the allocation function X must be non-decreasing in θ (Rochet, 1987). However, each consumer's expected utility function here does not satisfy the strict single-crossing property in (θ, X). In fact, as I show below in Section 3, X generally fails to be non-decreasing in θ.

The policymaker maximizes the expected total weighted surplus, consisting of: *(i)* consumer surplus; *(ii)* producer surplus; and *(iii)* her own profit. I proceed by describing each of these three components of the policymaker's objective function.

(i) **Consumer surplus.** Each consumer's surplus derives from his public option allocation and his private market outcome, and is weighted by a consumer-specific social welfare weight, denoted $\omega \in \mathbb{R}_+$. This represents the social value gained by giving that consumer one unit of money; throughout, I normalize $\mathbf{E}[\omega] = 1$. However, the policymaker does not observe each consumer's ω, but rather knows only the joint distribution of (θ, ω). Hence, she must infer each consumer's ω from his consumption behavior.

Each consumer's social welfare weight captures characteristics that determine how much the policymaker wishes to redistribute to him, but cannot be contracted on. The inability of the policymaker to contract on these characteristics captures the idea that the public option is "universal": often, the public option is required to guarantee access to all consumers by definition (Sitaraman and Alstott, 2019). Even when the universality requirement is relaxed, the policymaker might still be unable to contract on characteristics that determine ω such as race and religion due to legal or constitutional reasons. Therefore, the policymaker can choose direct mechanisms (X, T) that depend only on θ but not ω.

Even when these characteristics can potentially be contracted on, the policymaker might not be able to do so due to other reasons. One example is future health shocks, which neither the consumer nor the policymaker can foresee. Another example is future earnings potential, which could be the consumer's private information. While my formulation of the policymaker's problem might appear to preclude each consumer from reporting his ω (or any private information that is informative of ω), this assumption is without loss of generality according to standard arguments in the mechanism design literature (see, for example, Jehiel and Moldovanu, 2001; Che, Dessein, and Kartik, 2013a; and Dworczak ⓡ al., 2021): because ω does not affect consumer preferences, no direct revelation mechanism can condition allocations or payments directly on ω.

Each consumer is therefore assigned an expected weight of $\mathbf{E}[\omega \mid \theta]$ conditional on having a consumption type of θ. Given a policy (δ, X, T) that induces an aggregate quality Q, total weighted consumer surplus is given by

$$\begin{aligned} \mathrm{CS}(\delta, X, T; Q) = & \underbrace{\int_{\underline{\theta}}^{\overline{\theta}} \mathbf{E}[\omega \mid \theta] \left[u(\delta, \theta) X(\theta) - T(\theta)\right] \; \mathrm{d}F(\theta)}_{\text{expected weighted surplus from public consumption}} \\ & + \underbrace{\int_{\underline{\theta}}^{\overline{\theta}} \mathbf{E}[\omega \mid \theta] \cdot v_0(\theta, Q) \left[1 - X(\theta)\right] \; \mathrm{d}F(\theta)}_{\text{expected weighted surplus from private consumption}} . \end{aligned}$$

(ii) **Producer surplus.** Producer surplus arises because input supply is upward sloping, and is weighted equally across all producers with a welfare weight of $\alpha \in \mathbb{R}_+$. Since prices are equal to cost in the final good market, no producer surplus arises from conversion. Given a policy (δ, X, T) that induces an aggregate quality Q, weighted producer surplus is given by

$$\begin{aligned} \mathrm{PS}(Q) &= \alpha \Bigg[\underbrace{\int_{\underline{\theta}}^{\overline{\theta}} \left[\delta \cdot p(Q) X(\theta) + q_0(\theta, Q) \cdot p(Q) \left[1 - X(\theta)\right]\right] \; \mathrm{d}F(\theta)}_{\text{payment for input}} - \underbrace{\int_0^Q p(\hat{Q}) \; \mathrm{d}\hat{Q}}_{\text{cost of input}} \Bigg] \\ &= \alpha \left[Q \cdot p(Q) - \int_0^Q p(\hat{Q}) \; \mathrm{d}\hat{Q} \right] . \end{aligned}$$

(iii) **Policymaker's profit.** The policymaker's profit is equal to the difference between total payment and total cost for the public option, weighted by a welfare weight of 1. Because $\mathbf{E}[\omega] = 1$, the policymaker places the same weight on her own profit as average consumer surplus, which implies that uniform transfers between the policymaker and all consumers are welfare-neutral. Given a policy (δ, X, T) that induces an aggregate quality Q, the policymaker's profit is given by

$$\Pi(\delta, X, T; Q) = \underbrace{\int_{\underline{\theta}}^{\overline{\theta}} T(\theta) \; \mathrm{d}F(\theta)}_{\text{total payment}} - \underbrace{\int_{\underline{\theta}}^{\overline{\theta}} \left[c(\delta) + \kappa + \delta \cdot p(Q)\right] X(\theta) \; \mathrm{d}F(\theta)}_{\text{total cost}} .$$

The above discussion therefore allows the policymaker's objective function to be written as:

$$\mathrm{W}(\delta, X, T; Q) := \mathrm{CS}(\delta, X, T; Q) + \mathrm{PS}(Q) + \Pi(\delta, X, T; Q), \tag{OBJ}$$

where the equilibrium level of aggregate quality Q depends on the allocation function X via

$$Q = \underbrace{\int_{\underline{\theta}}^{\overline{\theta}} \delta X(\theta)\, \mathrm{d}F(\theta)}_{\text{aggregate public quality}} + \underbrace{\int_{\underline{\theta}}^{\overline{\theta}} q_0(\theta, Q)\, [1 - X(\theta)]\, \mathrm{d}F(\theta)}_{\text{aggregate private quality}}.$$

To conclude this section, I demonstrate that when $\alpha = 1$, the policymaker wishes to intervene only when θ and ω are correlated. This verifies the standard intuition (see, e.g., Aaron and Von Furstenberg, 1971) that in-kind redistribution is weakly dominated by cash transfers when consumers, producers, and the policymaker have the same welfare weight. The underlying reason is straightforward: each consumer exerts a pecuniary externality on other consumers through his consumption of quality. However, in the absence of redistributive motives, pecuniary externalities do not warrant intervention.

Proposition 2. *Suppose that θ and ω are independent and $\alpha = 1$. Then an optimal mechanism is $(X^*, T^*) = (0, 0)$, and in fact any optimal mechanism sets $X^* \equiv 0$ almost everywhere.*

3 Analysis

This section defines the direct and indirect effects of the public option and demonstrates how the policymaker's problem can be formulated in terms of their trade-off. I show how the policymaker's problem can be converted into a constrained mechanism design problem and characterize the resulting incentive constraints. Finally, I examine how these incentive constraints determine which consumers benefit from the public option.

3.1 Direct versus indirect effects

I begin by defining the direct and indirect effects of a policy change. To this end, consider the change in expected total weighted surplus due to a change in policy from (δ_0, X_0, T_0) to (δ, X, T). Let Q_0 and Q respectively denote the aggregate quality levels induced by (δ_0, X_0, T_0) and (δ, X, T) in equilibrium. Define the *direct effect* of this policy change as the change in expected total weighted surplus, holding fixed aggregate quality at the new level Q:

$$\Delta \mathrm{W}_D := \mathrm{W}(\delta, X, T; Q) - \mathrm{W}(\delta_0, X_0, T_0; Q).$$

Moreover, define the *indirect effect* of this policy change as the change in expected total weighted surplus due to the change in aggregate quality, evaluated at the initial policy (δ_0, X_0, T_0):

$$\Delta \mathrm{W}_I := \mathrm{W}(\delta_0, X_0, T_0; Q) - \mathrm{W}(\delta_0, X_0, T_0; Q_0).$$

Unless otherwise specified, I will evaluate the direct and indirect effects relative to the laissez-faire policy, $(\delta_0, X_0, T_0) = (0, 0, 0)$. Clearly, up to a constant term (independent of the chosen policy), the policymaker's objective (OBJ) is equal to the sum of the direct and indirect effects.

The direct and indirect effects each capture one of the two fundamental economic forces in the policymaker's problem. On one hand, the direct effect measures how well the policymaker can screen consumers; in particular, $\Delta \mathrm{W}_D = 0$ whenever $X = 0$. On the other hand, the indirect effect measures how much pecuniary externality the public option exerts on the private market; in particular, $\Delta \mathrm{W}_I = 0$ whenever input supply is perfectly elastic.

From a geometric perspective, the policymaker's problem can be viewed as a trade-off between the direct and indirect effects, as illustrated by Figure 1(a). The shaded region represents the set of outcomes that can be attained by feasible policies and randomization thereof, which includes the origin (as it is attained by the laissez-faire policy). The solid curve represents the Pareto frontier. The policymaker seeks to maximize the sum of the direct and indirect effects, which is equivalent to seeking the point of tangency on the Pareto frontier to linear indifference curves that slope downward at $-45°$.

This geometric approach shows that the ability to randomize between different policies does not benefit the policymaker. Indeed, if the Pareto frontier is not concave, as illustrated in Figure 1(b), the set of attainable outcomes would be the convex hull of points enclosed by the Pareto frontier. However, there must nonetheless be a point on both the Pareto frontier and its concavification that is tangent to a linear indifference curve that slopes downward at $-45°$. Consequently, restricting attention to deterministic policies entails no loss of generality for the purpose of solving for the optimal policy.

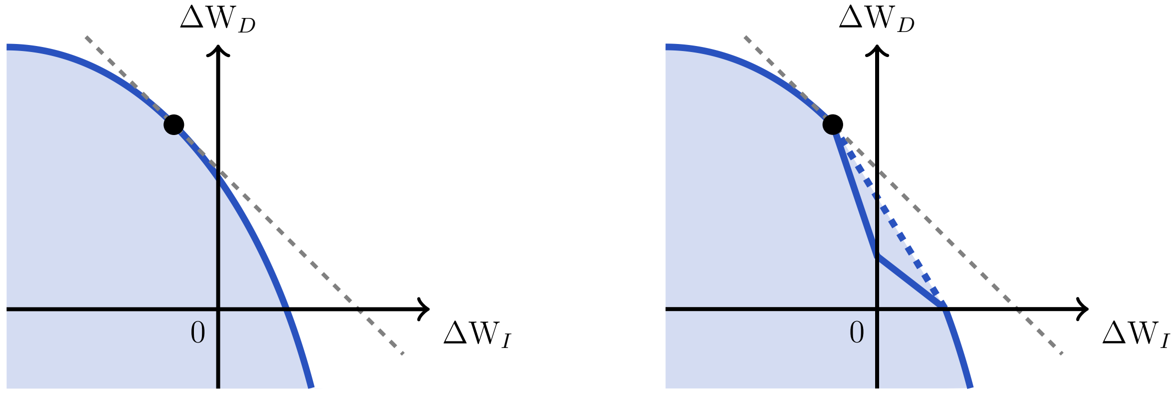

(a) *(b)*

Figure 1: Illustration of the trade-off between the direct and indirect effects.

3.2 Constrained mechanism design

Motivated by this geometric approach, I solve the policymaker's problem by characterizing policies that lie on the Pareto frontier. Heuristically, Figure 2 shows how this can be done by maximizing the direct effect for each value of the indirect effect.

Specifically, I solve the policymaker's problem in two stages. In the first stage, the policymaker chooses a quality level δ for the public option and an aggregate quality level Q that she wishes to induce in equilibrium. In the second stage, the policymaker chooses an incentive-compatible and individually rational mechanism (X, T), subject to the constraint that it induces the equilibrium aggregate quality level Q chosen in the first stage. The choice of Q uniquely pins down the value of the indirect effect; hence the policymaker's second-stage problem is equivalent to maximizing the direct effect for each value of the indirect effect for a given δ. I focus on the policymaker's second-stage problem for the remainder of this section.

The policymaker's second-stage problem is equivalent to a constrained mechanism design problem. Formally, the policymaker solves:

$$\max_{(X,T)} \; \mathrm{W}(\delta, X, T; Q) - \mathrm{W}(0, 0, 0; Q)$$

$$\text{s.t. } Q = \int_{\underline{\theta}}^{\overline{\theta}} \delta X(\theta) \, \mathrm{d}F(\theta) + \int_{\underline{\theta}}^{\overline{\theta}} q_0(\theta, Q) \left[1 - X(\theta)\right] \, \mathrm{d}F(\theta),$$

$$\theta \in \arg\max_{\theta' \in [\underline{\theta}, \overline{\theta}]} \left\{ u(\delta, \theta) \cdot X(\theta') + v_0(\theta, Q) \cdot [1 - X(\theta')] - T(\theta') \right\} \quad \text{for all } \theta \in [\underline{\theta}, \overline{\theta}], \tag{IC'}$$

$$u(\delta, \theta) \cdot X(\theta) + v_0(\theta, Q) \cdot [1 - X(\theta)] - T(\theta) \geq v_0(\theta, Q) \quad \text{for all } \theta \in [\underline{\theta}, \overline{\theta}]. \tag{IR'}$$

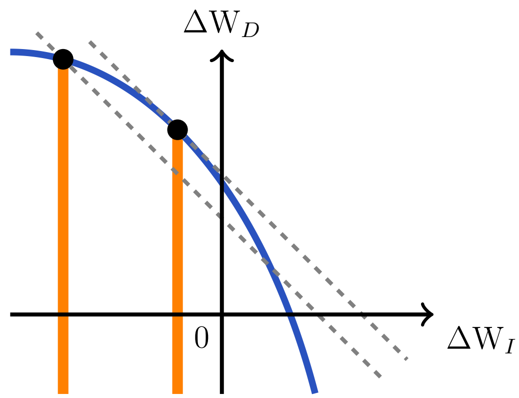

Figure 2: Solving the policymaker's problem by maximizing ΔW_D *for each value of* ΔW_I*.*

Unlike (IC) and (IR), notice that the aggregate quality Q in (IC') and (IR') no longer depends on the choices of δ and X. This comes at the expense of introducing a new constraint, namely, that the chosen mechanism must induce and equilibrium aggregate quality level of Q. I defer discussion of this new constraint to Section 4, where I derive the policymaker's optimal mechanism.

Next, I characterize the incentive constraints (IC') and (IR') in this constrained mechanism design problem by defining the effective consumption type. Because each consumer can always consume the private good, he does not value the public option at $u(\delta, \theta)$. Rather, his value of the public option is equal to the difference between $u(\delta, \theta)$ and what he would otherwise derive from private consumption, namely $v_0(\theta, Q)$. I refer to this as the *effective consumption type* of the consumer (with dependence on δ and Q suppressed for succinctness):

$$\eta(\theta) = u(\delta, \theta) - v_0(\theta, Q).$$

The effective consumption type η plays a crucial role in characterizing incentive constraints in the constrained mechanism problem. Its relevance can be inferred by rearranging (IC'), which yields:

$$\theta \in \arg\max_{\theta' \in [\underline{\theta}, \overline{\theta}]} \{\eta(\theta)X(\theta') - T(\theta')\} \quad \text{for all } \theta \in [\underline{\theta}, \overline{\theta}].$$

This is reminiscent of the incentive compatibility constraint in standard mechanism design models, where $\eta(\theta)$ replaces the role of the consumer's type. This analogy is reinforced by rearranging (IR'), which yields:

$$\eta(\theta)X(\theta) - T(\theta) \geq 0 \quad \text{for all } \theta \in [\underline{\theta}, \overline{\theta}].$$

The following lemma formalizes these observations.

Lemma 1. *For a given quality level $\delta \in \mathbb{R}_+$ of the public option, suppose that the policymaker's mechanism induces an equilibrium aggregate quality Q. Let*

$$\underline{\eta} = \min_{\theta \in [\underline{\theta}, \overline{\theta}]} \left[u(\delta, \theta) - v_0(\theta, Q) \right] \quad \text{and} \quad \overline{\eta} = \max_{\theta \in [\underline{\theta}, \overline{\theta}]} \left[u(\delta, \theta) - v_0(\theta, Q) \right].$$

Then any mechanism (X, T) satisfies (IC') and (IR') only if there exist a non-decreasing function $x : [\underline{\eta}, \overline{\eta}] \to [0, 1]$ and a function $t : [\underline{\eta}, \overline{\eta}] \to \mathbb{R}$ such that

(i) $X(\theta) = x(u(\delta, \theta) - v_0(\theta, Q))$ *almost everywhere; and*

(ii) $T(\theta) = t(u(\delta, \theta) - v_0(\theta, Q))$ *almost everywhere, such that*

$$\eta \cdot x(\eta) - t(\eta) = \underline{\eta} \cdot x(\underline{\eta}) - t(\underline{\eta}) + \int_{\underline{\eta}}^{\eta} x(s) \, \mathrm{d}s \quad \text{for all } \eta \in [\underline{\eta}, \overline{\eta}] \quad \text{and} \quad \underline{\eta} \cdot x(\underline{\eta}) - t(\underline{\eta}) \geq 0.$$

The intuition behind Lemma 1 is straightforward: given that consumption behavior for the public option depends only on consumers' effective consumption types, the policymaker cannot distinguish between two consumers who have the same effective consumption type. As such, she must give them the same allocation probabilities, even if they have different consumption types.

Lemma 1 demonstrates how the presence of a private market affects the incentive constraints that the policymaker faces. With a private market, the policymaker must now apply a change of variables: rather than screen consumers by their consumption types θ, she can screen consumers only by their effective consumption types η. Standard mechanism design methods apply once this change of variables is made. For the rest of this paper, I refer to (x, t) in Lemma 1 as the policymaker's *effective mechanism*, consisting of the *effective allocation function* $x : [\underline{\eta}, \overline{\eta}] \to [0, 1]$ and the *effective payment function* $t : [\underline{\eta}, \overline{\eta}] \to \mathbb{R}$.

Finally, an important implication of Lemma 1 is that allocation probability is quasiconcave in consumption type. Indeed, Lemma 1 demonstrates that allocation probability is non-decreasing in effective consumption type—unlike standard mechanism design models in which allocation probability is non-decreasing in consumption type. In turn, the effective consumption type $\eta(\theta)$ is quasiconcave in consumption type θ, as Figure 3 illustrates. The following proposition summarizes.

Proposition 3. *For any quality level $\delta \in \mathbb{R}_+$ of the public option and any mechanism (X, T) satisfying (IC'), the probability of receiving the public option $X(\theta)$ is quasiconcave in θ.*

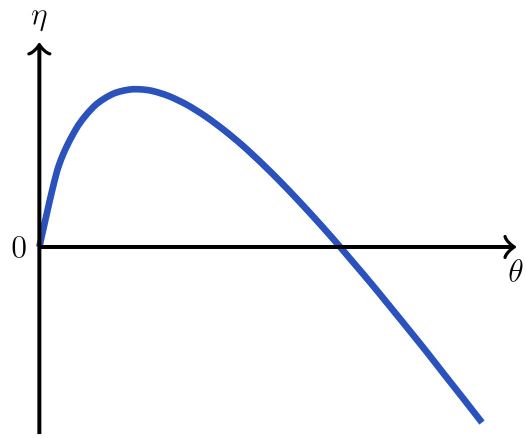

Figure 3: Effective consumption type η is quasiconcave in consumption type θ.

3.3 Impact of public option on consumer surplus

I conclude this section by demonstrating how the characterization of incentive constraints in the constrained mechanism design problem (Lemma 1) determines which consumers benefit from the public option.

To this end, I decompose the impact of public option on a given consumer's surplus into two components: the direct and indirect impacts on consumer surplus. Analogous to the direct effect as defined in Section 3.1, the *direct impact on consumer surplus* of a policy (δ, X, T) that induces an equilibrium aggregate quality level of Q is defined by

$$\Delta\mathrm{CS}_D(\theta) := \left[u(\delta,\theta) - v_0(\theta,Q)\right] X(\theta) - T(\theta).$$

This is the change in that consumer's surplus holding fixed the aggregate quality level at Q. Likewise, analogous to the indirect effect, the *indirect impact on consumer surplus* of a policy (δ, X, T) that induces an equilibrium aggregate quality level of Q is defined by

$$\Delta\mathrm{CS}_I(\theta) := v_0(\theta,Q) - v_0(\theta,Q_0),$$

where Q_0 denotes the laissez-faire equilibrium aggregate quality level. This is the change in that consumer's surplus due to the change in aggregate quality.

Proposition 4. *For any quality level $\delta \in \mathbb{R}_+$ and any mechanism (X, T) satisfying (IC'):*

(i) the direct impact on consumer surplus $\Delta\mathrm{CS}_D(\theta)$ is quasiconcave in θ; and

(ii) the indirect impact on consumer surplus $\Delta\mathrm{CS}_I(\theta)$ is non-increasing in θ if $Q > Q_0$ (and non-decreasing in θ if $Q < Q_0$).

On one hand, Proposition 4 shows that the direct impact on consumer surplus is generally non-monotone in consumption type. Indeed, the direct impact on consumer surplus is non-decreasing in effective consumption type by applying the usual envelope theorem argument to Lemma 1; in turn, effective consumption type is quasiconcave in consumption type (Figure 3). This shows how the well-known property of "no distortion at the top" in standard mechanism design models should be reinterpreted in the presence of a private market: there is no distortion for consumers with the highest effective consumption type—rather than those with the highest consumption type.

On the other hand, Proposition 4 also shows that the indirect impact on consumer surplus depends on whether aggregate quality increases or decreases. When the public option induces net upward substitution in quality on aggregate, then the prices of private goods increase, resulting in greater harm to consumers with higher consumption types. When the public option induces net downward substitution in quality on aggregate, then the prices of private goods decrease, resulting in greater benefit to consumers with higher consumption types.

Having established the relationship between impact on consumer surplus and consumption type, I now analyze the relationship between impact on consumer surplus and welfare weight. As discussed in Section 2, it is natural to suppose that consumption type and welfare weight are negatively correlated. One extreme is when $\theta = \phi(\omega)$ for some decreasing function $\phi : [\underline{\omega}, \overline{\omega}] \to [\underline{\theta}, \overline{\theta}]$, in which case a similar argument as that in Proposition 4 shows that the direct and indirect impacts on consumer surplus must also be quasiconcave and monotone respectively in welfare weight. Intuitively, this result must generalize as long as the correlation between consumption type and welfare weight is sufficiently strong, as the following result demonstrates.

Proposition 5. *Let $F_{\theta\,|\,\omega}$ denote the cumulative distribution function of $\theta\,|\,\omega$. Suppose that*

$$\theta\,|\,\omega_L \succeq_{\mathrm{FOSD}} \theta\,|\,\omega_H \quad \text{and} \quad \frac{\frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta\,|\,\omega_H)}{\frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta\,|\,\omega_L)} \quad \text{is increasing in } \theta \text{ for any } \underline{\omega} \leq \omega_L < \omega_H \leq \overline{\omega}.$$

Then, for any quality level $\delta \in \mathbb{R}_+$ and any mechanism (X, T) satisfying (IC'):

(i) the expected direct impact on consumer surplus $\mathbf{E}[\Delta\mathrm{CS}_D(\theta)\,|\,\omega]$ is quasiconcave in ω; and

(ii) the expected indirect impact on consumer surplus $\mathbf{E}[\Delta\mathrm{CS}_I(\theta)\,|\,\omega]$ is non-decreasing in ω if $Q > Q_0$ (and non-increasing in ω if $Q < Q_0$).

The first condition of Proposition 5 is relatively standard and captures the sense in which θ and ω are negatively correlated. It states that any decrease in the welfare weight shifts the distribution of θ in the sense of first-order stochastic dominance, so that consumers with lower ω tend to have higher θ. This condition is known to the statistics literature as "stochastic monotonicity" (see, e.g., Müller and Stoyan, 2002), and is a weaker condition than negative affiliation (e.g., Milgrom and Weber, 1982). This condition is commonly used in economics to describe correlation between random variables, including in recent papers by Haghpanah and Hartline (2021) and Yang (2021).

The second condition of Proposition 5 captures the sense in which the correlation between θ and ω is sufficiently strong. The condition is reminiscent of the monotone likelihood ratio property, with conditional density functions replaced instead by $\partial F_{\theta\,|\,\omega}/\partial\omega$. This derivative is well-defined because the first condition implies that $F_{\theta\,|\,\omega}(\theta\,|\,\cdot)$ is non-decreasing.

Jointly, the conditions of Proposition 5 are best understood through two examples:

(i) Suppose that $\theta = \phi(\omega)+\varepsilon$ for some decreasing function $\phi : [\underline{\omega}, \overline{\omega}] \to \mathbb{R}$ and a random variable ε with density function h and cumulative distribution function H, such that ω and ε are independent. Then

$$F_{\theta\,|\,\omega}(\theta\,|\,\omega) = \mathbf{P}[\phi(\omega) + \varepsilon \leq \theta] = H(\theta - \phi(\omega)) \implies \frac{\partial F_{\theta\,|\,\omega}}{\partial\omega}(\theta\,|\,\omega) = -\phi'(\omega)h(\theta - \phi(\omega)) \geq 0.$$

This verifies that the first condition is satisfied. Moreover, the second condition is satisfied if h is log-concave; in turn, many commonly used distributions have log-concave density functions, including the uniform, normal, and extreme-value distributions (see, e.g., Bagnoli and Bergstrom, 2005).

(ii) Suppose that $\theta = \varepsilon\phi(\omega)$ for some decreasing function $\phi : [\underline{\omega}, \overline{\omega}] \to \mathbb{R}$ and a positive random variable ε with density function h and cumulative distribution function H, such that ω and ε are independent. Then

$$F_{\theta\,|\,\omega}(\theta\,|\,\omega) = \mathbf{P}[\varepsilon\phi(\omega) \leq \theta] = H\left(\frac{\theta}{\phi(\omega)}\right) \implies \frac{\partial F_{\theta\,|\,\omega}}{\partial\omega}(\theta\,|\,\omega) = -\frac{\theta\phi'(\omega)}{[\phi(\omega)]^2} \cdot h\left(\frac{\theta}{\phi(\omega)}\right) \geq 0.$$

This verifies that the first condition is satisfied. Moreover, the second condition is satisfied if h is log-concave and non-increasing; common examples include the uniform and exponential distributions, and gamma distributions with shape parameter no greater than 1.

Under these conditions, Proposition 5 establishes the relationships between the expected direct and indirect impacts on consumer surplus and welfare weight. When each consumer's welfare weight represents his income or socioeconomic status, Proposition 5 is consistent with the empirical observation that that many public programs do not necessarily benefit poorer consumers more. This observation, commonly referred to as "Director's law," was first made by Aaron Director in the 1960s and has since attracted a number of political economy theories (Stigler, 1970). However, Proposition 5 shows that standard incentive constraints are consistent with Director's law. On one hand, consumers who derive the most expected surplus directly from the public option have intermediate values of welfare weights, and can hence be interpreted as having middle income or socioeconomic status. On the other hand, when the public option induces net upward aggregate quality substitution, consumers with lower welfare weights—who can be interpreted as having lower income or socioeconomic status—are disproportionately harmed by the increased private good prices. Here, incentive constraints, rather than political economy considerations, prevent the policymaker from giving the poor more expected surplus.

4 Optimal mechanisms

In this section, I show that the trade-off between the direct and indirect effects of the public option changes the qualitative nature of the optimal mechanism: it might require the policymaker to ration the public option.

4.1 Structure of the optimal mechanism

I begin by characterizing the structure of the optimal mechanism.

Theorem 1. *There exists an optimal mechanism* (X^*, T^*) *that is a menu of at most two prices, such that* $\operatorname{im} X^* \subseteq \{0, \pi, 1\}$ *for some* $0 < \pi < 1$.

Theorem 1 shows that the optimal mechanism might require rationing. Whereas consumers who pay the high price are allocated the public option with certainty, consumers who pay the low price are allocated with probability $\pi \in (0, 1)$.

Even though rationing might be required, Theorem 1 shows that the optimal mechanism nonetheless takes a simple form. A priori, the mechanism design problem is infinite-dimensional: the policymaker can choose any mechanism subject to (IC) and (IR), as characterized by Lemma 1. A posteriori, however, Theorem 1 shows that the problem reduces to a finite-dimensional one, consisting of at most three parameters: two prices and a rationing probability.

The proof of Theorem 1 reveals why the optimal mechanism takes such a simple form. Unlike recent papers (e.g., Dworczak ⓡ al., 2021; Loertscher and Muir, 2022) that derive similar results in their respective settings, I show that concavification methods do not generally apply here. Instead, I prove Theorem 1 by characterizing extreme points of all incentive-compatible mechanisms, which have a simple structure. This yields a non-constructive proof of Theorem 1, which I supplement in Appendix B by developing a method to explicitly compute the optimal mechanism.

In addition, the proof of Theorem 1 suggests that rationing might be optimal only because the policymaker must take into account the mechanism's effect on private good prices in equilibrium. In particular, standard arguments show that deterministic mechanisms are optimal in the absence of constraints (e.g., Harris and Raviv, 1981; Riley and Zeckhauser, 1983). By contrast, rationing becomes optimal only in the presence of constraints. While the mechanism design literature has considered constraints arising from capacity (e.g., Myerson, 1981; Bulow and Roberts, 1989) and budget (e.g., Samuelson, 1984), the policymaker in my setting does not face any such explicit constraints. Rather, she faces an implicit constraint that equilibrium prices in the private market depend on her choice of mechanism.

4.2 Proof of Theorem 1

I now provide a proof of Theorem 1 with technical details relegated to Appendix A. In particular, I defer the formal proof that the optimal mechanism exists and instead focus on its characterization.

The proof begins by solving the policymaker's constrained mechanism design problem given in Section 3.2. By Lemma 1, the policymaker equivalently chooses an effective mechanism (x, t); (IC') implies that x is increasing and t is obtained via the envelope theorem. Because the policymaker's welfare weight on her own profit is $\mathbf{E}[\omega] = 1$, uniform cash transfers between the policymaker and all consumers do not affect total weighted surplus. Thus the policymaker can normalize $t(\underline{\eta}) = 0$ without loss of generality, so that (IR') is satisfied.

Next, I characterize when the effective mechanism (x, t) induces an aggregate quality Q. Given Q, let G denote the cumulative distribution function of the effective consumption type η induced by the distribution F of consumption types, and let g denote the corresponding density function.

Lemma 2. *An effective allocation function x induces an aggregate quality Q in the market if and only if the following equilibrium condition is satisfied:*

$$Q = \int_{\underline{\eta}}^{\overline{\eta}} \left[\delta x(\eta) + \mathbf{E}[q_0(\theta, Q) \,|\, u(\delta, \theta) - v_0(\theta, Q) = \eta]\,[1 - x(\eta)]\right] \, \mathrm{d}G(\eta). \tag{E}$$

Lemma 2 is intuitive and proven formally in Appendix A: it states that the aggregate quality Q in equilibrium consists of the sum of quality from the public option and quality from the private good. The quality from the private good depends on Q via individual consumption $q_0(\theta, Q)$: an increase in Q raises private good prices, thereby reducing individual consumption. The proportion of public consumption to private consumption depends on the chosen allocation function. The equilibrium aggregate quality Q arises as a fixed point, as described by (E).

Given Lemmas 1 and 2 and holding fixed the aggregate quality Q, the policymaker's problem can be written as follows, omitting terms that do not explicitly depend on x:

$$\max_{x} \int_{\underline{\eta}}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta \cdot p(Q) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \mid s] - 1] \;\mathrm{d}G(s)}{g(\eta)}\right] x(\eta) \;\mathrm{d}G(\eta) \tag{P}$$

$$\text{s.t.} \begin{cases} x : [\underline{\eta}, \overline{\eta}] \to [0,1] \text{ is non-decreasing,} \\ Q - \mathbf{E}[q_0(\theta, Q)] = \displaystyle\int_{\underline{\eta}}^{\overline{\eta}} [\delta - \mathbf{E}[q_0(\theta, Q) \mid \eta]] \, x(\eta) \;\mathrm{d}G(\eta). \end{cases}$$

The policymaker's problem (P) is reminiscent of the classical monopoly pricing problem, except with the usual quantity constraint replaced by the equilibrium constraint (E). Recall that the quantity constraint in the classical monopoly pricing problem arises because the monopolist has a limited quantity of the good to sell. By contrast, the equilibrium constraint (E) arises because the policymaker must induce a pre-determined level of aggregate quality Q in the market.

Unlike the classical monopoly pricing problem, however, the policymaker's problem (P) cannot generally be solved with concavification methods. This is because the usual quantity constraint always tightens with each additional unit of good sold. As such, the shadow cost of each additional unit of good is always positive. By contrast, the equilibrium constraint (E) does not always tighten with each additional unit of the public option allocated. Instead, when the policymaker allocates the public option to a consumer with consumption type η, whether the constraint tightens or slackens depends on the sign of $\delta - \mathbf{E}[q_0(\theta, Q) \mid \eta]$. As such, the Lagrange multiplier associated with the constraint has ambiguous sign in general. Of course, the Lagrange multiplier might be signed when additional assumptions are introduced, such as when the public option induces downward substitution (cf. Definition 1). Under these assumptions, concavification methods can be used to solve the policymaker's problem via a constructive approach, as I show below in Proposition 8.

Instead, I solve the policymaker's problem (P) using linear programming methods. To formalize this, I denote the extreme points of any set S by $\operatorname{ex} S$ and use the following theorem from infinite-dimensional concave programming based on the work of Bauer (1958) and Szapiel (1975):

Theorem. *Let K be a convex, compact set in a locally convex Hausdorff space, and let $\ell : K \to \mathbb{R}^m$ be a continuous affine function such that $\Sigma \subseteq \operatorname{im} \ell$ is a closed and convex set. Suppose that $\ell^{-1}(\Sigma)$ is nonempty and and that $\Omega : K \to \mathbb{R}$ is a continuous convex function. Then there exists $z^* \in \ell^{-1}(\Sigma)$ such that $\Omega(z^*) = \max_{z \in \ell^{-1}(\Sigma)} \Omega(z)$ and*

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

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

In the present context, K represents the space of implementable effective allocation functions viewed as a subset of $L^1([\underline{\eta}, \overline{\eta}])$, defined by

$$K := \left\{ x : [\underline{\eta}, \overline{\eta}] \to [0, 1] \text{ is non-decreasing} \right\}.$$

Recall that two functions in L^1 are equal if they agree almost everywhere on their domain; hence the above definition of K simply requires that x is almost everywhere equal to a non-decreasing function. As I show in Appendix A, K is convex, compact (in the L^1 topology), and is a subset of a normed linear space. It is also well-known that the set of extreme points of K consists of the step functions (cf. Skreta, 2006 and Manelli and Vincent, 2007):

Lemma 3. *The function $x \in L^1([\underline{\eta}, \overline{\eta}])$ is an extreme point of K if and only if x is a non-decreasing function satisfying* $\operatorname{im} x \subseteq \{0, 1\}$.

Lemma 3 shows that the candidate solutions to the policymaker's problem in the absence of the equilibrium constraint (E) are exactly those implementable by a single price. In particular, when input supply is perfectly elastic, the policymaker no longer can affect aggregate quality Q through her choice of allocation function X. As such, she no longer has to include the equilibrium constraint (E) in her problem, so $\operatorname{im} X^* \subset \{0, 1\}$ for her optimal allocation function X^*.

To complete the proof of Theorem 1, observe that the policymaker's objective function in (P) is linear (hence convex) and continuous in x, so that the results of Bauer (1958) and Szapiel (1975) apply. Define the function $\ell : K \to \mathbb{R}$ and the set $\Sigma \subseteq \mathbb{R}$ by

$$\begin{cases} \ell(x) & := \int_{\underline{\eta}}^{\overline{\eta}} [\delta - \mathbf{E}[q_0(\theta, Q) \mid \eta]]\, x(\eta)\, \mathrm{d}G(\eta), \\ \Sigma & := \{Q - \mathbf{E}[q_0(\theta, Q)]\}. \end{cases}$$

It is easy to verify that ℓ is continuous and linear, and that Σ is closed and convex. Moreover, whenever $\ell^{-1}(\Sigma)$ is nonempty, the optimal effective allocation function x^* can be written as the convex combination of at most two extreme points of K. By Lemma 3, this implies the existence of $\pi \in (0, 1)$ such that $\operatorname{im} x^* \subseteq \{0, \pi, 1\}$. This completes the proof of Theorem 1.

4.3 When is rationing optimal?

While Theorem 1 shows that the optimal mechanism might require rationing, I now demonstrate how the optimality of rationing arises from the trade-off between the direct and indirect effects of the public option.

To this end, I begin by showing that no rationing is required in the absence of a trade-off between the direct and indirect effects—that is, when the allocation of the public option cannot affect prices of private goods.

Proposition 6. *When input supply is perfectly elastic, there exists an optimal mechanism* (X^*, T^*) *that does not require rationing, such that* $\operatorname{im} X^* \subseteq \{0, 1\}$.

Proposition 6 thus shows that the trade-off between the direct and indirect effects is necessary for the optimal mechanism to require rationing. To understand why, suppose that the policymaker optimally rations some consumers when the supply of input is perfectly elastic. As the direct effect of the public option is linear in allocation probability, the policymaker can always weakly increase the direct effect by allocating to the same set of consumers with either higher or lower probability. Moreover, because input supply is perfectly elastic, the indirect effect remains the same (i.e., zero) and incentive constraints are not affected as the price of input (and hence the prices of private goods) remains the same. As such, it would also be optimal for the policymaker to allocate the public option with probability 0 or 1 to these consumers, thereby removing any need for rationing in the optimal mechanism.

The trade-off between the direct and indirect effects can be mathematically formalized via the Lagrange multiplier λ on the equilibrium constraint (E) in the policymaker's problem (P). The policymaker's Lagrangian can thus be written as

$$L(x;\lambda,Q)=\lambda\left[Q-\mathbf{E}[q_0(\theta,Q)]+\int_{\underline{\eta}}^{\overline{\eta}}\left[\mathbf{E}[q_0(\theta,Q)\,|\,\eta]-\delta\right]x(\eta)\ \mathrm{d}G(\eta)\right]$$
$$+\int_{\underline{\eta}}^{\overline{\eta}}\left[\eta-c(\delta)-\kappa-\delta p(Q)+\frac{\int_{\eta}^{\overline{\eta}}\left[\mathbf{E}[\omega\,|\,s]-1\right]\ \mathrm{d}G(s)}{g(\eta)}\right]x(\eta)\ \mathrm{d}G(\eta).$$

When input supply is perfectly elastic, there is no trade-off between the direct and indirect effects, which corresponds to the case $\lambda=0$. Then the resulting Lagrangian can be easily maximized over the set of implementable effective allocation functions, $K:=\{x:[\underline{\eta},\overline{\eta}]\to[0,1]\text{ is non-decreasing}\}$. This problem is thus analogous to the classical monopoly pricing problem with a constant marginal cost, for which no rationing is required (see, e.g., Wilson, 1988 and Bulow and Roberts, 1989). In general, however, $\lambda\in\mathbb{R}$ can be either positive or negative and represents the shadow price of ensuring that the the aggregate demand for quality is equal to its aggregate supply Q.

Next, I demonstrate how the trade-off between the direct and indirect effects can result in the optimal rationing of the public option. The Lagrangian approach suggests two cases to examine: $\lambda>0$ and $\lambda<0$.

On one hand, when $\lambda>0$, then rationing can help to either enhance the indirect effect (when positive) or mitigate the indirect effect (when negative) when the mechanism induces consumers to substitute to goods of higher quality on average. In this case, the binding inequality is

$$\int_{\underline{\eta}}^{\overline{\eta}}\left[\delta-\mathbf{E}[q_0(\theta,Q)\,|\,\eta]\right]x(\eta)\ \mathrm{d}G(\eta)\le Q-\mathbf{E}[q_0(\theta,Q)].$$

If $Q>Q_0$, then the right-hand side is positive; thus, when the mechanism induces consumers to substitute to goods of higher quality on average, the policymaker wishes to restrict the increase in quality consumption that the public option induces. Similarly, if $Q<Q_0$, then the right-hand side is negative; thus, when the mechanism induces consumers to substitute to goods of higher quality on average, the policymaker wishes to expand the decrease in quality consumption that the public option induces.

To illustrate this intuition, consider an example in which the policymaker places a much higher welfare weight on consumers with the lowest consumption types, so that the direct effect is large and positive only when these consumers are allocated the public option. Incentive constraints compel the policymaker to allocate the public option to some consumers with higher consumption types (and low welfare weights) in order to allocate it to these consumers.

However, if the policymaker does not ration, then this might raise the price of input significantly if this policy induces consumers to substitute to goods of higher quality on average. In turn, this could result in a large and negative indirect effect as the policymaker is hurt by a much higher cost of supplying the public option. By contrast, rationing could limit the negative indirect effect—especially when the supply of input is initially relatively elastic but grows increasingly inelastic—while still allowing a fraction of the large and positive direct effect to be realized.

On the other hand, when $\lambda < 0$, then a similar logic arises when the mechanism induces consumers to substitute to goods of lower quality on average. In this case, the binding inequality is

$$\int_{\underline{\eta}}^{\overline{\eta}} \left[\mathbf{E}[q_0(\theta, Q) \,|\, \eta] - \delta\right] x(\eta) \, \mathrm{d}G(\eta) \leq \mathbf{E}[q_0(\theta, Q)] - Q.$$

If $Q > Q_0$, then the right-hand side is negative; thus, when the mechanism induces consumers to substitute to goods of lower quality on average, the policymaker wishes to expand the decrease in quality consumption that the public option induces. Likewise, if $Q < Q_0$, then the right-hand side is positive; thus, when the mechanism induces consumers to substitute to goods of lower quality on average, the policymaker wishes to restrict the increase in quality consumption that the public option induces.

This discussion motivates the question of when the mechanism induces consumers to substitute to goods of higher or lower quality on average. Intuitively, rationing arises when the equilibrium constraint (E) binds and the coefficient of $x(\eta)$ in the integrand in the Lagrangian fails to be non-decreasing in η. This depends on the signs of both λ and the derivative of $\mathbf{E}[q_0(\theta, Q) \,|\, \eta]$: rationing is "likelier" to be required when either *(i)* $\lambda > 0$ and $\mathbf{E}[q_0(\theta, Q)]$ is decreasing or *(ii)* $\lambda < 0$ and $\mathbf{E}[q_0(\theta, Q)]$ is increasing. To formalize this intuition, consider the auxiliary problem ($\mathrm{P}_\varnothing$) in which the policymaker is constrained to use a mechanism with no rationing:

$$\max_{Q \in \mathbb{R}_+} \left[\mathbf{E}[\omega v_0(\theta, Q)] + \alpha \, \mathrm{PS}(Q) + \min_{\lambda \in \mathbb{R}} \max_{x \in K_\varnothing} L(x; \lambda, Q)\right], \tag{$\mathrm{P}_\varnothing$}$$

where

$$K_\varnothing := \left\{x : [\underline{\eta}, \overline{\eta}] \to \{0, 1\} \text{ is non-decreasing}\right\}.$$

Proposition 7. *For any solution $(x_\varnothing, \lambda_\varnothing, Q_\varnothing)$ to the auxiliary problem ($\mathrm{P}_\varnothing$), let the corresponding quantity of the public option allocated be denoted by*

$$m_\varnothing := \int_{\underline{\eta}}^{\overline{\eta}} x_\varnothing(\eta) \, \mathrm{d}G(\eta),$$

and suppose that

$$\frac{\mathrm{d}}{\mathrm{d}\eta}\left[\eta+\frac{\int_{\eta}^{\overline{\eta}}\left[\mathbf{E}[\omega\,|\,\eta=s]-1\right]\,\mathrm{d}G(s)}{g(\eta)}+\lambda_{\varnothing}\,\mathbf{E}[q_0(\theta,Q_{\varnothing})\,|\,\eta]\right]\Bigg|_{\eta=G^{-1}(1-m_{\varnothing})}<0.$$

Then any optimal mechanism (X^*, T^*) *requires rationing; that is, there exists* $\pi \in (0,1)$ *such that* $\pi \in \operatorname{im} X^*$*.*

By showing that the sign of the derivative of $\mathbf{E}[q_0(\theta,Q)\,|\,\eta]$ determines whether the mechanism induces consumers to substitute to goods of higher or lower quality on average, Proposition 7 provides a sufficient condition for the optimal mechanism to require rationing based on the trade-off between the direct and indirect effects.

While the derivative of $\mathbf{E}[q_0(\theta,Q)\,|\,\eta]$ is difficult to sign in general, I now show how a regularity condition can be imposed to help construct the optimal mechanism—and hence characterize when rationing is optimal.

Definition 1. *The public option* induces downward substitution *if*

$$\delta<\mathbf{E}\left[q_0(\theta,Q)\,|\,u(\delta,Q)-v_0(\theta,Q)=\eta\right]\quad \textit{for any } Q\geq 0 \textit{ and } \underline{\eta}\leq\eta<\overline{\eta}.$$

Definition 1 requires the expected quality difference between the public option and private good to have the same sign across all effective consumption types at a given aggregate quality Q. The public option induces downward substitution if the expected quality from private consumption for each effective consumption type is higher than that from the public option. Consequently, each consumption type substitutes to a lower quality level in expectation when he applies for the public option. The public option thus exerts downward pressure on the prices of private goods at Q.

Definition 1 is intuitively related to how the expected quality of private consumption varies with effective consumption type. In particular, if the expected quality of private consumption decreases with effective consumption type, then the following condition guarantees that the public option induces downward substitution:

$$u(\delta,\overline{\theta})-u(\delta,\underline{\theta})\leq\lim_{Q\to+\infty}\left[v_0(\overline{\theta},Q)-v_0(\underline{\theta},Q)\right]\quad\text{and}\quad\delta>q_0(\underline{\theta},0).\tag{I}$$

This condition captures the intuition that δ is "interior" in the sense that: *(i)* δ is sufficiently small so that the public option is more attractive to consumers with low consumption types; and *(ii)* δ is nonetheless bigger than what these consumers would otherwise consume in the private market. Under condition (I), the highest effective consumption type $\overline{\eta}$ consumes a private good of quality exactly equal to δ if he is not allocated the public option, so

$$\delta = \mathbf{E}\left[q_0(\theta, Q) \,|\, u(\delta, Q) - v_0(\theta, Q) = \overline{\eta}\right] \quad \text{for any } Q \geq 0.$$

Since the expected quality of private consumption decreases with effective consumption type, this implies that the public option induces downward substitution.

Definition 1 is not vacuous. For example, simulations show that Definition 1 is satisfied with a uniform distribution $F(\theta) = \theta$ for $\theta \in [0, 1]$, a utility function $u(q, \theta) = \sqrt{\theta q}$, and an input supply curve $p(Q)$ that satisfies $p(Q) \leq 1/(4\sqrt{\delta})$ for all $Q \in \mathbb{R}_+$. While it is possible to define a similar condition for the public option to induce upward substitution, under condition (I), the public option cannot induce upward substitution except if $u(\delta, \underline{\theta}) - v_0(\underline{\theta}, Q) = \underline{\eta} = u(\delta, \overline{\theta}) - v_0(\overline{\theta}, Q)$ for all $Q \in \mathbb{R}_+$. Otherwise, condition (I) implies that effective consumption types in a neighborhood of $\underline{\eta}$ consume private goods of quality strictly higher than δ.

In turn, the ratio of densities of high consumption types to low consumption types determines how the expected quality of private consumption varies with effective consumption type. To illustrate this, consider the case where input supply is perfectly elastic. Then when the distribution of consumption types is uniform and the quality of the public option δ satisfies condition (I), it can be shown that the public option induces downward substitution. However, the public option might not induce downward substitution when more density is shifted from high consumption types to low consumption types.

Definition 1 yields a sufficient condition for rationing via a constructive characterization of the optimal mechanism, as opposed to the non-constructive characterization given in Theorem 1.

Proposition 8. *Suppose that the public option induces downward substitution, and suppose that there exists an optimal mechanism (X^*, T^*) that induces an aggregate quality Q^*. Define*

$$\begin{cases} H(\eta) := \dfrac{\int_{\underline{\eta}}^{\eta} \left[\mathbf{E}[q_0(\theta, Q^*) \,|\, \eta = s] - \delta\right] \, \mathrm{d}G(s)}{\mathbf{E}[q_0(\theta, Q^*)] - \delta}, \\ \Psi(r) := \displaystyle\int_{H^{-1}(1-r)}^{\overline{\eta}} \frac{\eta - c(\delta) - \kappa - \delta p(Q^*) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \,|\, \eta = s] - 1] \, \mathrm{d}G(s)}{g(\eta)}}{\mathbf{E}[q_0(\theta, Q^*) \,|\, \eta] - \delta} \, \mathrm{d}H(\eta). \end{cases}$$

For any function ϕ, let $\operatorname{co} \phi$ denote the concave closure of ϕ (i.e., the pointwise smallest concave *function that bounds ϕ from above). Then $\pi \in \operatorname{im} X^*$ for some $\pi \in (0,1)$ if and only if*

$$\Psi\left(\frac{\mathbf{E}[q_0(\theta, Q^*)] - Q^*}{\mathbf{E}[q_0(\theta, Q^*)] - \delta}\right) \neq \operatorname{co}\Psi\left(\frac{\mathbf{E}[q_0(\theta, Q^*)] - Q^*}{\mathbf{E}[q_0(\theta, Q^*)] - \delta}\right).$$

Proposition 8 shows that, if the public option induces downward substitution, the optimality of rationing depends on two functions, namely, H and Ψ. The former, $H(\eta)$, denotes the fraction of reduction in aggregate quality consumption that arises from allocating the public option to all effective consumption types below η. The latter, $\Psi(r)$, is the cumulative sum of ratios across consumers with effective consumption types $\eta \geq H^{-1}(1-r)$ of the expected marginal gain in social welfare arising from the public option to the expected marginal reduction in quality consumption. That is, the weight assigned to the ratio for each effective consumption type η is not the density $g(\eta)$ of that effective consumption type, but rather the density weighted by that effective consumption type's contribution to aggregate quality reduction $h(\eta) = H'(\eta)$.

Proposition 8 demonstrates that the optimal mechanism can require rationing if Ψ fails to be concave—that is, if the ratio of the expected marginal gain in social welfare from allocating the public option to an consumer with effective consumption type η to the expected marginal reduction in his quality consumption is decreasing. Indeed, there are two reasons why the policymaker might wish to allocate the public option to a consumer. First, there might be a high expected marginal gain in social welfare from allocating the public option to that consumer; this effect is captured by the numerator of the ratio. Second, allocating to that consumer might help reduce aggregate quality consumption to a larger extent, which exerts greater downward pressure on private market prices; this effect is captured by the denominator of the ratio.

The intuition behind this ratio can be understood by viewing the policymaker's problem as a continuous knapsack problem. Using an incentive-compatible mechanism, the policymaker packs as many effective consumption types as possible into her knapsack. The size of her knapsack is determined by the amount of reduction in quality consumption that she wishes to induce. In turn, the size of each effective consumption type is determined by how much he reduces his quality consumption if allocated the public option. Meanwhile, the value of each effective consumption type is determined by how much he contributes to the gain in social welfare if allocated the public option. To solve this continuous knapsack problem, the policymaker computes the ratio of value to size—that is, the "bang for the buck"—for each effective consumption type.

Proposition 8 also connects the policymaker's problem to similar problems considered by the literature (e.g., Dworczak ⓡ al., 2021; Loertscher and Muir, 2022). In particular, it shows how characterization methods in the literature can be applied to the policymaker's problem, but only under the additional assumption that the public option induces downward substitution. Indeed, a similar approach fails in general because allocating to an consumer might either increase or decrease aggregate quality consumption in expectation. This necessitates a different approach to characterize the optimal mechanism, as given in the proof of Theorem 1.

In turn, Proposition 8 provides a modified sufficient condition for rationing to be required in the spirit of Proposition 7:

Corollary 1. *Suppose that the public option induces downward substitution, and suppose that any solution* $(x_\varnothing, \lambda_\varnothing, Q_\varnothing)$ *to the auxiliary problem* ($P_\varnothing$) *satisfies* $\delta < Q_\varnothing < Q_0$ *and*

$$\frac{\mathrm{d}}{\mathrm{d}\eta}\left[\frac{\eta - c(\delta) - \kappa - \delta p(Q_\varnothing) + \frac{\int_\eta^{\overline{\eta}}[\mathbf{E}[\omega \mid s]-1]\ \mathrm{d}G(s)}{g(\eta)}}{\mathbf{E}[q_0(\theta, Q_\varnothing) \mid \eta] - \delta}\right]\Bigg|_{\eta = H^{-1}\left(\frac{Q_\varnothing - \delta}{\mathbf{E}[q_0(\theta, Q_\varnothing)] - \delta}\right)} < 0$$

where the effective consumption type is $\eta = u(\delta, \theta) - v_0(\theta, Q_\varnothing)$*. Then any optimal mechanism* (X^*, T^*) *requires rationing; that is, there exists* $\pi \in (0,1)$ *such that* $\pi \in \operatorname{im} X^*$*.*

Unlike Proposition 7, Corollary 1 provides a sufficient condition that depends only on the total quantity of public option allocated, $Q_\varnothing$. This is because the assumption that the public option induces downward substitution allows the Lagrange multiplier on the equilibrium constraint (E) to be signed; hence the sufficient condition for rationing to be optimal can be expressed independently of the Lagrange multiplier.

Finally, a different sufficient condition for the optimal mechanism to require rationing can be obtained in terms of the quantity of the public option allocated.

Proposition 9. *Suppose that* $q_0(\underline{\theta}, Q_0) < \delta < q_0(\overline{\theta}, Q_0)$ *and* $\kappa > 0$*. Then there exists* $m \in (0,1)$ *such that, if every solution* $(x_\varnothing, \lambda_\varnothing, Q_\varnothing)$ *to the auxiliary problem* ($P_\varnothing$) *satisfies*

$$0 < m_\varnothing := \int_{\underline{\eta}}^{\overline{\eta}} x_\varnothing(\eta)\ \mathrm{d}G(\eta) < m,$$

then any optimal mechanism (X^*, T^*) *requires rationing; that is, there exists* $\pi \in (0,1)$ *such that* $\pi \in \operatorname{im} X^*$*.*

Proposition 9 shows that rationing is optimal when the quantity of the public option allocated is sufficiently small. Indeed, if a small quantity of the public option is allocated without rationing, then the policymaker incurs a loss from crowding out the private market since the policymaker is less efficient. Moreover, as the policymaker clears the market for the public option, its price must be sufficiently high; thus consumers who receive the public option are approximately indifferent between the public option and the private good that they would otherwise have consumed. Hence, if an optimal mechanism allocates a small quantity of the public option, then it requires rationing.

5 Optimal quality

In this section, I show the trade-off between the direct and indirect effects changes how the quality level δ of the public option should be chosen: it introduces a novel incentive in the policymaker's first-order condition for quality.

5.1 Marginal effect of quality

The main result of this section is the following characterization of the marginal effect of quality on total weighted surplus.

Theorem 2. *For any policy* (δ, X, T) *that induces an aggregate quality* Q *such that* (X, T) *satisfies* (IC'), *the marginal effect of quality on total weighted surplus is*

$$\frac{\partial \mathrm{W}}{\partial \delta} = \underbrace{\int_{\underline{\theta}}^{\overline{\theta}} \left[\frac{\partial u}{\partial q}(\delta, \theta) - c'(\delta) - p(Q)\right] X(\theta)\ \mathrm{d}F(\theta)}_{\textit{difference between marginal benefit and marginal cost}}$$

$$\underbrace{- \int_{\underline{\theta}}^{\overline{\theta}} \frac{\int_{\theta}^{\overline{\theta}} [1 - \mathbf{E}[\omega \mid s]]\ \mathrm{d}F(s)}{f(\theta)} \cdot \frac{\partial^2 u}{\partial q\, \partial \theta}(\delta, \theta) X(\theta)\ \mathrm{d}F(\theta)}_{\textit{screening distortion}}$$

$$\underbrace{+ \left[\alpha Q - \int_{\underline{\theta}}^{\overline{\theta}} [\delta X(\theta) + \mathbf{E}[\omega \mid \theta] q_0(\theta, Q)\, [1 - X(\theta)]]\ \mathrm{d}F(\theta)\right] \cdot \frac{\partial P}{\partial \delta}}_{\textit{pecuniary externality distortion}},$$

where $\partial P / \partial \delta$ *denotes the marginal effect of quality on the price of input, given by*

$$\frac{\partial P}{\partial \delta} = \frac{p'(Q) \int_{\underline{\theta}}^{\overline{\theta}} X(\theta)\ \mathrm{d}F(\theta)}{1 - \int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial q_0}{\partial Q}(\theta, Q)\, [1 - X(\theta)]\ \mathrm{d}F(\theta)}.$$

Theorem 2 shows that the marginal effect of quality on total weighted surplus consists of three terms: the difference between marginal benefit and marginal cost, the screening distortion, and the pecuniary distortion. Unless the optimal quality is to set a baseline level of quality (i.e., $\delta^* = 0$), the sum of these three terms must be equal to zero at the optimal policy (δ^*, X^*, T^*).

The first term—the difference between marginal benefit and marginal cost—is standard and describes the marginal effect of quality on total weighted surplus when the policymaker has a utilitarian objective function. This is obtained as a special case of Theorem 1 when $\alpha = \mathbf{E}[\omega \mid \theta] \equiv 1$, in which case both the screening and pecuniary externality distortions vanish.

The second term—the screening distortion—arises from the policymaker's screening problem: she wants to discourage those with higher consumption types from mimicking the consumption behavior of those with lower consumption types. In turn, this screening problem arises because the policymaker has a redistributive objective function and consumers have private information about their own consumption types. The screening distortion is negative when expected welfare weight $\mathbf{E}[\omega \mid \theta]$ is decreasing in consumption type θ (so that those with lower consumption types are poorer in expectation). This captures the well-known intuition that the policymaker has an incentive to distort quality downwards in order to better target poorer consumers (e.g., Nichols and Zeckhauser, 1982; Besley and Coate, 1991; Gahvari and Mattos, 2007).

The third term—the pecuniary externality distortion—is new to my model and arises from the pecuniary externality that the public option exerts on the private market. In the special case where input supply is perfectly elastic, there is no pecuniary distortion. In general, however, an increase in the quality level of the public option results in a first-order increase in aggregate quality consumption in the market—and hence a first-order increase in the price for input:

$$\frac{\partial P}{\partial \delta} = \frac{p'(Q) \int_{\underline{\theta}}^{\overline{\theta}} X(\theta) \ \mathrm{d}F(\theta)}{1 - \int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial q_0}{\partial Q}(\theta, Q) \left[1 - X(\theta)\right] \ \mathrm{d}F(\theta)} \geq 0.$$

The sign of the pecuniary externality distortion depends on the difference between weighted aggregate quality supplied and weighted aggregate quality demanded. On one hand, the weighted aggregate quality supplied is equal to the aggregate quality level Q, weighted by the welfare weight α on private producers. On the other hand, the weighted aggregate quality demanded is the sum of aggregate quality demanded via the public option $\mathbf{E}[\delta X(\theta)]$, weighted by the policymaker's welfare weight of 1, and the quality demanded by each consumer in the private market $q_0(\theta, Q) \left[1 - X(\theta)\right]$, weighted by his (expected) welfare weight of $\mathbf{E}[\omega \mid \theta]$. When $\alpha = 0$, the weighted aggregate quality supplied vanishes, in which case the pecuniary externality distortion is non-positive and creates an incentive for the policymaker to further reduce the quality of the public option.

Finally, the magnitude of the pecuniary externality distortion depends not only on weighted aggregate quality supplied and weighted aggregate quality demanded, but is also larger when each consumer's individual demand for quality and the aggregate supply for quality are more inelastic. This is easiest to see when $c(q) \equiv 0$, in which case the marginal effect of quality on the price of input can be expressed in terms of each consumer's individual demand elasticity for quality $\varepsilon_D(\cdot, \theta)$ and the aggregate supply elasticity for quality $\varepsilon_S(\cdot)$:

$$\frac{\partial P}{\partial \delta} = \frac{p(Q) \int_{\underline{\theta}}^{\overline{\theta}} X(\theta) \, \mathrm{d}F(\theta)}{\varepsilon_S(p(Q)) \cdot Q - \int_{\underline{\theta}}^{\overline{\theta}} \varepsilon_D(p(Q), \theta) \cdot q_0(\theta, Q) \left[1 - X(\theta)\right] \, \mathrm{d}F(\theta)}.$$

Similar, albeit more complicated, expressions can be derived for the general case when $c(q) \not\equiv 0$. Therefore, the pecuniary externality distortion is larger for goods with more inelastic demand for quality, which are commonly interpreted as necessities, and goods with more inelastic supply for input, such as space for housing.

5.2 When is baseline quality optimal?

I conclude this section by examining when it is optimal for the policymaker to choose a baseline level of quality, namely, $\delta^* = 0$.

To this end, I impose an additional restriction on consumer utility, namely, that $u(q, \theta) = \theta \nu(q)$ for some twice continuously differentiable function ν. This additional restriction is made only for simplicity and is widely imposed in screening models in the literature.

First, I guve a sufficient condition under which a baseline quality is not optimal.

Proposition 10. *Suppose that the policymaker optimally supplies a public option, so that $X^* \not\equiv 0$. Then the optimal quality satisfies $\delta^* \neq 0$ when*

$$\theta - \frac{\int_{\theta}^{\overline{\theta}} \left[1 - \mathbf{E}[\omega \,|\, s]\right] \, \mathrm{d}F(s)}{f(\theta)} \geq 0 \qquad \textit{for all } \theta \in [\underline{\theta}, \overline{\theta}],$$

with the inequality holding strictly for a subset of consumption types that has positive measure.

The sufficient condition given in Proposition 10 holds for many distributions of consumption types and welfare weights. First, observe that the condition holds when $\mathbf{E}[\omega \,|\, \theta] = 1$ for any distribution of consumption types; and, by continuity, it must hold when there is sufficiently small dispersion in consumer welfare weights (i.e., when inequality is sufficiently low in the market).

Second, observe that the condition is satisfied for any distribution of welfare weights under the stronger condition that

$$\theta - \frac{1 - F(\theta)}{f(\theta)} > 0 \qquad \text{for all } \theta \in [\underline{\theta}, \overline{\theta}].$$

This latter condition is satisfied, for example, when consumption types have a Pareto distribution with finite mean. Finally, the sufficient condition given in Proposition 10 is also satisfied when the expression is strictly quasiconcave in θ, an assumption that has been shown to be relatively permissive in the literature (cf. Assumption 1 of Dworczak ⓡ al., 2021).

Under this sufficient condition, the policymaker does not optimally provide the lowest possible quality. Intuitively, the Inada condition that consumer utility functions satisfy (cf. Assumption 1) has two implications: while consumers derive a high marginal utility for quality when quality is low, a low-quality public option also helps the policymaker screen more effectively. Under the condition in Proposition 10, the former dominates, so the policymaker does not provide the lowest possible quality.

Next, I give a sufficient condition under which a baseline quality is optimal.

Proposition 11. *Suppose that the policymaker optimally supplies a public option, so that $X^* \not\equiv 0$. Then the optimal quality satisfies $\delta^* = 0$ when $\alpha = 0$ and*

$$\int_{\underline{\theta}}^{\overline{\theta}} \left[\theta - \frac{\int_{\theta}^{\overline{\theta}} \left[1 - \mathbf{E}[\omega \mid s]\right] \, \mathrm{d}F(s)}{f(\theta)} \right] X^*(\theta) \, \mathrm{d}F(\theta) \le 0.$$

While the sufficient condition given in Proposition 11 is strong, it leads to the stark conclusion that the policymaker provides the lowest possible quality. The assumption that $\alpha = 0$ means that the policymaker has an incentive not only to redistribute from rich consumers to poor consumers, but also from producers to consumers. This assumption ensures that the pecuniary externality distortion (cf. Theorem 2) is negative. Together with the latter assumption on the distributions of consumption types and welfare weights, this guarantees that the marginal effect of quality on total weighted surplus must be negative at the lowest possible quality level.

Under this sufficient condition, the policymaker provides the lowest possible quality because the benefit in screening effectiveness outweighs the harm to consumers. Such a low-quality public option attracts only the poorest consumers, whom the policymaker pays (i.e., $T^* < 0$). These payments function as "conditional cash transfers" that compensate consumers for consuming at a lower quality than they otherwise would have in the private market, in the spirit of Gahvari and Mattos (2007).

Finally, the optimality of baseline quality can arise only when the policymaker is restricted to non-negative levels of quality. In reality, the policymaker might be able to design "ordeals" as in Nichols and Zeckhauser (1982), which can be accommodated by allowing for negative levels of quality. This would generally allow the policymaker to obtain a higher total weighted surplus, and the optimal quality level would be characterized by setting the marginal effect of quality on total weighted surplus to zero as described in the discussion following Theorem 2.

6 Extensions

In this section, I show how my analysis generalizes in three different dimensions that have featured in policy discussions about the public option. First, I show how externalities and paternalism can be incorporated into the policymaker's objective. Second, I present the implications of market power on my results. Third, I discuss how my analysis extends when the policymaker has access to additional policy instruments.

6.1 Externalities and paternalism

I begin by incorporating externalities and paternalism into the policymaker's objective function. This is motivated by the fact that externalities and paternalism are commonly cited reasons for the use of public options (Currie and Gahvari, 2008). Moreover, many public options (e.g., public schools and health care) are used in markets where consumption produces externalities.

I model externalities and paternalistic preferences by adding a term that depends on weighted aggregate consumption to the policymaker's objective function. Formally, in addition to θ and ω, each consumer is now also endowed with a potentially heterogeneous $\xi \in [\underline{\xi}, \overline{\xi}] \subseteq \mathbb{R}_+$ that captures how much externality he produces per unit of quality consumed. As with θ and ω, ξ is modeled as private information (e.g., mental health externalities due to living in a bigger apartment), which can be alternatively interpreted as a limitation of what information the policymaker can use to screen consumers. Consequently, if each consumer with consumption type θ consumes a good of quality $q(\theta)$, aggregate externality is given by

$$E = \int_{\underline{\theta}}^{\overline{\theta}} \mathbf{E}[\xi \mid \theta] q(\theta) \, \mathrm{d}F(\theta).$$

The policymaker incorporates this into her objective via an additively separable term $e(E)$, where $e : \mathbb{R}_+ \to \mathbb{R}$ is continuously differentiable and referred to as the externality function.

This way of modeling externalities and paternalistic preferences allows me to flexibly capture a wide range of situations. For example, a paternalistic policymaker might directly value aggregate quality consumption Q; this is achieved by setting $\xi \equiv 1$ for all consumers. More generally, the weight ξ that the policymaker places on each consumer's quality consumption might differ between consumers with the same θ and ω. Finally, the restriction that the aggregate externality enters the policymaker's objective function is made only for the sake of simplicity; more general forms of externality can also be analyzed (Kang, 2022).

The inclusion of externalities and paternalistic preferences potentially changes how the policymaker should allocate a public option, as the following result indicates.

Theorem 3. *There exists an optimal mechanism* (X^*, T^*) *that is a menu of at most three prices, such that* $\operatorname{im} X^* \subset \{0, \pi_1, \pi_2, 1\}$ *for some* $0 < \pi_1 < \pi_2 < 1$*. Moreover, if the externality function* e *is convex, then at most two prices are required.*

Theorem 3 shows that externalities and paternalism provide an additional reason to ration the public option. In contrast to Theorem 1, a menu of two prices might no longer be sufficient for the optimal mechanism; instead, three prices might be required. In particular, Theorem 3 implies that rationing the public option might be necessary even when input supply is perfectly elastic, in contrast to Proposition 6. This is because rationing might give lower consumption types a higher allocation probability, which increases the policymaker's objective function when ξ and θ are negatively correlated. However, as Theorem 3 shows, this argument depends on the curvature of the externalities function. Specifically, if the externality function is convex, then the policymaker would no longer wish to ration due to the increasing marginal benefit that she realizes from increasing the aggregate externality. In that case, the interaction between the public option and private market remains the sole reason to ration the public option, as in Theorem 1.

With the inclusion of externalities and paternalistic preferences also changes, the sign of the indirect effect due to externalities depends on whether the externality function is increasing or decreasing at E_0. This is intuitive: when e is increasing, the externality can be interpreted as a positive externality. Raising the aggregate externality level then raises the policymaker's objective function, which indicates that the public option is more valuable. Conversely, the public option is less valuable when e is decreasing, in which case the externality can be interpreted as a negative externality.

6.2 Imperfect competition

Next, I present the implications of imperfect competition on my results. This is motivated by the fact that imperfect competition in private markets has been cited as a reason for the use of public options (Sitaraman and Alstott, 2019); moreover, many markets in which governments use a public option are imperfectly competitive (e.g., health care, food and nutrition).

I begin by considering the effects of imperfect competition among private producers who convert input into the good. Then $c(q)$ should be reinterpreted as the price of a private good of quality q, rather than the cost of a private good of quality q. As long as Assumption 2 continues to hold, each consumer's consumption problem admits a unique solution. Consequently, the analysis changes only in two ways. First, producer surplus must now include the profit that private producers make from converting input into the good. Second, it is now possible for the policymaker's inefficiency parameter κ to be negative if she is more efficient than imperfectly competitive private producers. As long as the welfare weight α on private producers is sufficiently small, imperfect competition in the conversion market gives the policymaker more reason to use a public option.

Next, I consider the effects of imperfect competition in the market for the input good. To this end, suppose for simplicity that $c(q) \equiv 0$ and that input is sold by a monopolist; other imperfectly competitive market structures can be similarly accommodated (e.g., Mahoney and Weyl, 2017). Let $C(Q)$ denote the cost for Q units of input, and suppose that C is increasing and convex, so that marginal cost is increasing. Then my analysis extends under the following additional assumption:

Assumption 3. *In addition to the conditions specified in Assumption 1, the utility function u is homogeneous of degree 1 in (q, θ) such that the consumer with the lowest consumption type $\underline{\theta}$ has log-concave demand for quality; that is,*

$$p \mapsto \left(\frac{\partial u}{\partial q}\right)^{-1}(p; \underline{\theta}) \quad \textit{is log-concave.}$$

Assumption 3 is best illustrated with a parametric example. Consider $u(q, \theta) = \theta^{1/\varepsilon} q^{1-1/\varepsilon}$ for $\varepsilon > 0$, which is clearly homogeneous of degree 1. In this case, it can be readily verified that each consumer's demand for quality has a constant elasticity of ε:

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

Then the second condition in Assumption 3 is also satisfied as log-demand is concave in p:

$$\log D(p, \underline{\theta}) = \log \underline{\theta} + \varepsilon \left[\log \left(\frac{\varepsilon - 1}{\varepsilon} \right) - \log p \right].$$

More generally, Assumption 3 allows a similar characterization of the equilibrium condition to that given in Lemma 2. Indeed, Assumption 3 guarantees that there is a one-to-one mapping between the residual demand curve and residual marginal revenue curve for quality. Equilibrium is obtained at the unique point of intersection between the residual marginal revenue curve and marginal cost curve, rather than at the point of intersection between the residual demand curve and marginal cost curve under perfect competition. By contrast, in the absence of Assumption 3, the policymaker can separately affect the residual demand curve and residual marginal revenue curve, thereby complicating the analysis.

In turn, Assumption 3 allows a similar characterization of optimal mechanisms to that given in Theorem 1, except that an additional rationing option is required. The need for an additional rationing option arises because the policymaker must account for not just the residual demand at the input price that she wishes to induce, but also the elasticity of residual demand. Nevertheless, this shows that the intuition from Theorem 1 continues to hold even under imperfect competition: that is, rationing might be required when the public option can affect private market prices.

Theorem 4. *There exists an optimal mechanism (X^*, T^*) that is a menu of at most three prices, such that* $\operatorname{im} X^* \subset \{0, \pi_1, \pi_2, 1\}$ *for some* $0 < \pi_1 < \pi_2 < 1$.

Finally, imperfect competition in the input market changes the indirect effect. This difference arises not just because imperfect competition changes the producer surplus, but also because the equilibrium condition changes. Consequently, the set of mechanisms that effect a given change in aggregate quality ΔQ changes. Unlike the case of imperfect competition in the conversion market, it is thus no longer clear whether imperfect competition gives the policymaker more reason to use a public option.

6.3 Additional policy instruments

Finally, I discuss how my analysis extends when the policymaker has access to additional policy instruments, namely when she can impose uniform tax or subsidy on quality in the market.

In this case, my analysis holds for any level of tax or subsidy that the policymaker may wish to impose. Indeed, the policymaker's problem can now be decomposed into three stages: the policymaker first chooses the optimal level of tax or subsidy on quality, then chooses the level of aggregate quality that she wishes to induce, and finally chooses a mechanism that induces this level of aggregate quality. My analysis holds in the second and third stages of this decomposition, and in particular must hold for the optimal level of tax or subsidy that the policymaker chooses in the first stage. As such, similar characterizations of incentive constraints (cf. Lemma 1) and optimal mechanisms (cf. Theorem 1) can be obtained.

It is important to note that the value of a public option can nonetheless be positive even if the policymaker can impose a uniform tax or subsidy on quality. This is because the policymaker can generally still improve on the optimal uniform tax or quality, as I show in a companion paper (Kang, 2022). There, I solve for the optimal nonlinear tax or subsidy that the policymaker might wish to impose when she cannot further redistribute with a public option. Following the logic of Atkinson and Stiglitz (1976), I conjecture that a public option is no longer optimal once the policymaker can impose a nonlinear tax or subsidy on quality. In practice, however, a nonlinear tax or subsidy might be difficult to implement when input can be costlessly resold.

7 Conclusion

When governments provide a public option, the prices of private goods are often affected; yet this fundamental economic force is often left out in analyses that assume that the policymaker can design the entire market. In this paper, I develop a tractable model in which the policymaker can design only part of the market, which I use to study the equilibrium effects of a public option on the private market and how they impact optimal design.

I find that these equilibrium effects qualitative change the nature of optimal mechanisms and the optimal choice of quality. These equilibrium effects provide a new justification for rationing the public option, which allows the policymaker to vary residual demand for quality and residual supply of input in the private market. In so doing, the policymaker can either enhance or mitigate the effect of the public option on the prices of private goods. This might help explain the widespread use of rationing in the allocation of public options, such as lotteries in public housing and waiting times in public health care. These equilibrium effects also create new incentives to raise or lower the quality of the public option, which depends on the relative welfare weights that the policymaker assigns to consumers and producers.

My results provide a framework to understand optimal redistribution in different markets where the magnitude of these equilibrium effects might differ. These equilibrium effects tend to be larger for markets for which demand and supply for quality tend to be more inelastic, which can be interpreted as markets for "necessities." These findings complement the results of the literature on public provision, which tend to focus on direct effects and the ability to target consumers by exploiting the statistical correlation between welfare weight and consumption type.

While I have employed a long-run analysis in this paper by allowing the distribution of quality to be endogenously determined, similar intuitions are likely to hold in the short to medium run too. The distribution of quality supplied is likely to be more constrained in the short to medium run, which creates an additional incentive for the policymaker to ration. In addition, input supply is likely to be more inelastic in the short to medium run than in the long run, which suggests more scope for redistribution with a public option.

Finally, this paper offers avenues for future research. While I have focused on the case of unit demand for the good, the case of continuous demand entails the issue of whether the policymaker should allow consumption of the public option to be supplemented with consumption of the private good, which Currie and Gahvari (2008) refer to as "topping up." Furthermore, the framework provided in this paper could also be used to empirically estimate the value of a public option in different markets. A comparison between a public option and existing policies would enhance our understanding of when the public option should be used in practice.

References

Aaron, H. J. and G. M. Von Furstenberg (1971): "The Inefficiency of Transfers in Kind: The Case of Housing Assistance," *Western Economic Journal*, 9, 184–191.

Akbarpour, M. ⓡ E. Budish ⓡ P. Dworczak ⓡ S. D. Kominers (2021): "An Economic Framework for Vaccine Prioritization," *Working paper*.

Akbarpour, M. ⓡ P. Dworczak ⓡ S. D. Kominers (2022): "Redistributive Allocation Mechanisms," *Working paper*.

Atal, J. P., J. I. Cuesta, F. González, and C. Otero (2021): "The Economics of the Public Option: Evidence from Local Pharmaceutical Markets," *Working paper*.

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

Aumann, R. J. and M. Maschler (1995): *Repeated Games with Incomplete Information*, Cambridge, MA: MIT Press.

Bagnoli, M. and T. Bergstrom (2005): "Log-Concave Probability and Its Applications," *Economic Theory*, 26, 445–469.

Barron, E. N. (1983): "Remarks on a Simple Optimal Control Problem with Monotone Control Functions," *Journal of Optimization Theory and Applications*, 41, 573–586.

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

Baum-Snow, N. and J. Marion (2009): "The Effects of Low Income Housing Tax Credit Developments on Neighborhoods," *Journal of Public Economics*, 93, 654–666.

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.

Bulow, J. and J. Roberts (1989): "The Simple Economics of Optimal Auctions," *Journal of Political Economy*, 97, 1060–1090.

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.

Che, Y.-K., W. Dessein, and N. Kartik (2013a): "Pandering to Persuade," *American Economic Review*, 103, 47–79.

Che, Y.-K., I. Gale, and J. Kim (2013b): "Assigning Resources to Budget-Constrained Agents," *Review of Economic Studies*, 80, 73–107.

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

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

Curto, V., L. Einav, A. Finkelstein, J. Levin, and J. Bhattacharya (2019): "Health Care Spending and Utilization in Public and Private Medicare," *American Economic Journal: Applied Economics*, 11, 302–32.

Diamond, R. and T. McQuade (2019): "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development," *Journal of Political Economy*, 127, 1063–1117.

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

Dinerstein, M., C. Neilson, and S. Otero (2020): "The Equilibrium Effects of Public Provision in Education Markets: Evidence from a Public School Expansion Policy," *Working paper.*

Dinerstein, M. and T. D. Smith (2021): "Quantifying the Supply Response of Private Schools to Public Policies," *American Economic Review*, 111, 3376–3417.

Doval, L. and V. Skreta (2022): "Constrained Information Design," *Working paper.*

Duggan, M. and F. M. Scott Morton (2006): "The Distortionary Effects of Government Procurement: Evidence from Medicaid Prescription Drug Purchasing," *Quarterly Journal of Economics*, 121, 1–30.

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.

Epple, D. and R. E. Romano (1998): "Competition between Private and Public Schools, Vouchers, and Peer-Group Effects," *American Economic Review*, 88, 33–62.

Friedman, M. (1962): *Capitalism and Freedom*, Chicago, IL: University of Chicago Press.

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.

Haghpanah, N. and J. Hartline (2021): "When Is Pure Bundling Optimal?" *Review of Economic Studies*, 88, 1127–1156.

Handbury, J. and S. Moshary (2021): "School Food Policy Affects Everyone: Retail Responses to the National School Lunch Program," *Working paper*.

Harris, M. and A. Raviv (1981): "A Theory of Monopoly Pricing Schemes with Demand Uncertainty," *American Economic Review*, 71, 347–365.

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

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

Hoxby, C. M. (2000): "Does Competition among Public Schools Benefit Students and Taxpayers?" *American Economic Review*, 90, 1209–1238.

Jehiel, P. and B. Moldovanu (2001): "Efficient Design with Interdependent Valuations," *Econometrica*, 69, 1237–1259.

Jiménez-Hernández, D. and E. Seira (2021): "Should the Government Sell You Goods? Evidence from the Milk Market in Mexico," *Working paper.*

Kamenica, E. and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

Kang, M. and C. Z. Zheng (2020): "Pareto Optimality of Allocating the Bad," *Working paper.*

Kang, Z. Y. (2022): "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.*

Kleiner, A., B. Moldovanu, and P. Strack (2021): "Extreme Points and Majorization: Economic Applications," *Econometrica*, 89, 1557–1593.

Lasserre, J. B. (2009): *Moments, Positive Polynomials and Their Applications*, London, UK: Imperial College Press.

Le Treust, M. and T. Tomala (2019): "Persuasion with Limited Communication Capacity," *Journal of Economic Theory*, 184, 104940.

Loertscher, S. and E. V. Muir (2022): "Monopoly Pricing, Optimal Rationing, and Resale," *Journal of Political Economy*, 130, 566–635.

Mahoney, N. and E. G. Weyl (2017): "Imperfect Competition in Selection Markets," *Review of Economics and Statistics*, 99, 637–651.

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

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

Milgrom, P. R. and R. J. Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122.

Milgrom, P. and I. Segal (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70, 583–601.

Müller, A. and D. Stoyan (2002): *Comparison Methods for Stochastic Models and Risks* Volume 389, West Sussex, UK: John Wiley & Sons.

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

Nichols, A. L. and R. J. Zeckhauser (1982): "Targeting Transfers through Restrictions on Recipients," *American Economic Review*, 72, 372–377.

Nikzad, A. (2022): "Constrained Majorization: Applications in Mechanism Design," *Working paper.*

Pai, M. and P. Strack (2022): "Taxing Externalities Without Hurting the Poor," *Working paper.*

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

Piazzesi, M. and M. Schneider (2016): "Housing and Macroeconomics," in *Handbook of Macroeconomics* ed. by Taylor, J. B. and Uhlig, H. Volume 2: Elsevier, Chap. 19, 1547–1640.

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.*

Riley, J. and R. Zeckhauser (1983): "Optimal Selling Strategies: When to Haggle, When to Hold Firm," *Quarterly Journal of Economics*, 98, 267–289.

Rochet, J.-C. (1987): "A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context," *Journal of Mathematical Economics*, 16, 191–200.

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

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

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

Stigler, G. J. (1970): "Director's Law of Public Income Redistribution," *Journal of Law and Economics*, 13, 1–10.

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

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

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

Waldinger, D. (2021): "Targeting In-Kind Transfers Through Market Design: A Revealed Preference Analysis of Public Housing Allocation," *American Economic Review*, 111, 2660–96.

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.

Wilson, C. A. (1988): "On the Optimal Pricing Policy of a Monopolist," *Journal of Political Economy*, 96, 164–176.

Yang, F. (2021): "Costly Multidimensional Screening," *Working paper.*

Appendix A Omitted proofs

A.1 Proofs from Section 2

A.1.1 Proof of Proposition 1

Given the level of aggregate quality Q, consumers solve the utility maximization problem

$$\max_{q \in \mathbb{R}_+} \left[u(q, \theta) - c(q) - q \cdot p(Q) \right].$$

For any consumer with consumption $\theta > 0$, observe that an interior solution exists because

$$\lim_{q \to 0} \frac{\partial u}{\partial q}(q, \theta) = +\infty > c'(0) \quad \text{and} \quad \lim_{q \to +\infty} \frac{\partial u}{\partial q}(q, \theta) = 0 < p(Q) \leq \lim_{q \to +\infty} c'(q) + p(Q).$$

Consequently, any solution q must be characterized by the first-order condition

$$\frac{\partial u}{\partial q}(q, \theta) - c'(q) = p(Q).$$

The solution is unique because u is strictly concave and c is convex, so that the left-hand side of the above equation is decreasing in q and the right-hand side is constant in q.

Let $q_0(\theta, Q)$ denote the unique solution. Observe that $q_0(\theta, Q)$ is non-increasing in Q: in the above equation, the right-hand side is non-decreasing in Q while the left-hand side is decreasing in q. Moreover, because u is strictly concave, $\partial u / \partial q$ is decreasing; hence the inverse function of $\partial u / \partial q$ exists. Denote this by $D(\cdot, \theta)$. Therefore

$$q_0(\theta, Q) = D\left(c'(q_0(\theta, Q)) + p(Q), \theta\right).$$

In particular, this must hold for the equilibrium level of aggregate quality Q_0, which must satisfy the fixed-point condition

$$Q = \int_{\underline{\theta}}^{\overline{\theta}} q_0(\theta, Q) \, \mathrm{d}F(\theta).$$

Finally, observe that Q_0 is uniquely defined since the left-hand side of the above equation is increasing in Q while the right-hand side is non-increasing in Q since each $q_0(\theta, \cdot)$ is non-increasing.

A.1.2 Proof of Proposition 2

Given that θ and ω are independent and $\alpha = 1$, the policymaker's objective function (OBJ) reduces to the standard utilitarian objective. Because the policymaker is weakly less efficient than private producers and because the private market for the good and the input market are perfectly competitive, hence the laissez-faire equilibrium is first-best.

A.2 Proofs from Section 3

A.2.1 Proof of Lemma 1

Conditional on the level of aggregate quality being $Q(X) = X$, (IC') can be written as

$$\theta \in \arg\max_{\theta' \in [\underline{\theta}, \overline{\theta}]} \{[u(\delta, \theta) - v_0(\theta, Q)]\, X(\theta) - T(\theta)\} \quad \text{for all } \theta \in [\underline{\theta}, \overline{\theta}].$$

Hence, for any $\theta, \theta' \in [\underline{\theta}, \overline{\theta}]$,

$$u(\delta, \theta) - v_0(\theta, Q) > u(\delta, \theta') - v_0(\theta', Q) \implies X(\theta) \geq X(\theta'). \tag{A1}$$

Let $\eta = u(\delta, \theta) - v_0(\theta, Q)$. By the envelope theorem,

$$\frac{\partial \eta}{\partial \theta} = \frac{\partial u}{\partial \theta}(\delta, \theta) - \frac{\partial u}{\partial \theta}(q_0(\theta, Q), \theta).$$

Since u satisfies the strict single-crossing property in (q, θ) by Assumption 1, it follows that

$$\frac{\partial \eta}{\partial \theta} \begin{cases} > 0 & \text{for } q_0(\theta, Q) < \delta, \\ < 0 & \text{for } q_0(\theta, Q) > \delta. \end{cases}$$

Since $q_0(\cdot, Q)$ is increasing (cf. Proposition 1), it follows that η is strictly quasiconcave in θ. Let $\theta^* \in \arg\max_{\theta \in [\underline{\theta}, \overline{\theta}]} [u(\delta, \theta) - v_0(\theta, Q)]$, so that for each $\eta \in [\underline{\eta}, \overline{\eta}]$ there exist at most two values of θ, namely $\theta_H(\eta) > \theta^*$ and $\theta_L(\eta) \leq \theta^*$, such that

$$u(\delta, \theta_L(\eta)) - v_0(\theta_L(\eta), Q) = \eta = u(\delta, \theta_H(\eta)) - v_0(\theta_H(\eta), Q).$$

Define

$$y(\eta,\theta) := \begin{cases} X(\theta_H(\eta)) & \text{if } \theta > \theta^*, \\ X(\theta_L(\eta)) & \text{if } \theta \leq \theta^*. \end{cases}$$

For sufficiently small $\varepsilon > 0$, equation (A1) implies that

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

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

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

In turn, this implies that $y(\eta,\theta) = x(\eta)$ for almost every $\eta \in [\underline{\eta},\overline{\eta}]$. Thus $X(\theta_H(\eta)) = X(\theta_L(\eta)) = x(\eta)$, where x is a non-decreasing function, for almost every $\eta \in [\underline{\eta},\overline{\eta}]$. This proves statement *(i)* of the lemma; statement *(ii)* follows from the envelope theorem of Milgrom and Segal (2002).

A.2.2 Proof of Proposition 3

This follows from the proof of Lemma 1: since $\eta(\theta)$ is quasiconcave in θ and $x(\eta)$ is non-decreasing in η, it follows that $X(\theta) = x \circ \eta(\theta)$ must be quasiconcave in θ.

A.2.3 Proof of Proposition 4

(i) This follows from the proof of Lemma 1: By the envelope theorem,

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

Thus $v(\eta)$ is non-decreasing in η. Since $\eta(\theta)$ is quasiconcave in θ and $v(\eta)$ is non-decreasing in η, it follows that $\Delta\mathrm{CS}_D(\theta) = v \circ \eta(\theta)$ must be quasiconcave in θ.

(ii) By the envelope theorem,

$$\begin{aligned} \frac{\partial\,\Delta\mathrm{CS}_I}{\partial\theta}(\theta) &= \frac{\partial v_0}{\partial\theta}(\theta,Q) - \frac{\partial v_0}{\partial\theta}(\theta,Q_0) \\ &= \frac{\partial u}{\partial\theta}(q_0(\theta,Q),\theta) - \frac{\partial u}{\partial\theta}(q_0(\theta,Q_0),\theta). \end{aligned}$$

Since u satisfies the strict single-crossing property in (q, θ) by Assumption 1, it follows that

$$\frac{\partial\, \Delta\mathrm{CS}_I}{\partial \theta}(\theta) \begin{cases} > 0 & \text{for } q_0(\theta, Q) > q_0(\theta, Q_0), \\ < 0 & \text{for } q_0(\theta, Q) < q_0(\theta, Q_0). \end{cases}$$

Since $q_0(\theta, \cdot)$ is non-increasing (cf. Proposition 1), it follows that $\Delta\mathrm{CS}_I(\theta)$ is non-increasing in θ if $Q > Q_0$ (and non-decreasing in θ if $Q < Q_0$).

A.2.4 Proof of Proposition 5

(i) Since $\theta \,|\, \omega_L \succeq_{\mathrm{FOSD}} \theta \,|\, \omega_H$, $F_{\theta\,|\,\omega}(\theta \,|\, \cdot)$ is non-decreasing. Hence the derivative $\partial F_{\theta\,|\,\omega}/\partial \omega$ is defined and non-negative almost everywhere. By Proposition 4, there exists $\theta^* \in [\underline{\theta}, \overline{\theta}]$ such that $\partial\, \Delta\mathrm{CS}_D(\theta)/\partial\theta \geq 0$ for $\theta < \theta^*$ and $\partial\, \Delta\mathrm{CS}_D(\theta)/\partial\theta \leq 0$ for $\theta > \theta^*$. Observe that each consumer's expected surplus from the public option conditional on his welfare weight ω is

$$\begin{aligned} \mathbf{E}\left[\Delta\mathrm{CS}_D(\theta, Q) \,|\, \omega\right] &= \int_{\underline{\theta}}^{\overline{\theta}} \Delta\mathrm{CS}_D(\theta)\ \mathrm{d}F_{\theta\,|\,\omega}(\theta \,|\, \omega) \\ &= \Delta\mathrm{CS}_D(\overline{\theta}) - \int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial\, \Delta\mathrm{CS}_D}{\partial \theta}(\theta) F_{\theta\,|\,\omega}(\theta \,|\, \omega)\ \mathrm{d}\theta. \end{aligned}$$

Consequently,

$$\frac{\partial}{\partial \omega} \mathbf{E}\left[\Delta\mathrm{CS}_D(\theta) \,|\, \omega\right] = -\int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial\, \Delta\mathrm{CS}_D}{\partial \theta}(\theta) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta \,|\, \omega)\ \mathrm{d}\theta.$$

To show that $\omega \mapsto \mathbf{E}\left[\Delta\mathrm{CS}_D(\theta) \,|\, \omega\right]$ is quasiconcave, it suffices to show that

$$\frac{\partial}{\partial \omega} \mathbf{E}\left[\Delta\mathrm{CS}_D(\theta) \,|\, \omega\right]\Big|_{\omega=\omega_L} \leq 0 \implies \frac{\partial}{\partial \omega} \mathbf{E}\left[\Delta\mathrm{CS}_D(\theta) \,|\, \omega\right]\Big|_{\omega=\omega_H} \leq 0 \quad \text{for } \underline{\omega} \leq \omega_L < \omega_H \leq \overline{\omega}.$$

To this end, define $\psi : [\underline{\theta}, \theta^*] \to [\theta^*, \overline{\theta}]$ implicitly by

$$\int_{\theta}^{\theta^*} \frac{\partial\, \Delta\mathrm{CS}_D}{\partial \theta}(\theta) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta \,|\, \omega_L)\ \mathrm{d}\theta + \int_{\theta^*}^{\psi(\theta)} \frac{\partial\, \Delta\mathrm{CS}_D}{\partial \theta}(\theta) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta \,|\, \omega_L)\ \mathrm{d}\theta = 0.$$

Note that ψ is well-defined because

$$\frac{\partial}{\partial \omega} \mathbf{E}\left[\Delta \mathrm{CS}_D(\theta) \,|\, \omega\right]\Big|_{\omega=\omega_L} \leq 0 \quad \text{by assumption.}$$

Implicit differentiation yields

$$-\frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\theta) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta \,|\, \omega_L) + \frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\psi(\theta)) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\psi(\theta) \,|\, \omega_L) \psi'(\theta) = 0.$$

Therefore

$$\begin{aligned} 0 &= -\frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\theta) \cdot \frac{\frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta \,|\, \omega_L)}{\frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\psi(\theta) \,|\, \omega_L)} + \frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\psi(\theta))\psi'(\theta) \\ &\leq -\frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\theta) \cdot \frac{\frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta \,|\, \omega_H)}{\frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\psi(\theta) \,|\, \omega_H)} + \frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\psi(\theta))\psi'(\theta). \end{aligned}$$

This implies that

$$0 \geq -\frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\theta) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta \,|\, \omega_H) + \frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\psi(\theta)) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\psi(\theta) \,|\, \omega_H) \psi'(\theta).$$

Using the fact that $\psi(\theta^*) = \theta^*$, integration yields

$$\begin{aligned} 0 &\leq \int_{\underline{\theta}}^{\theta^*} \left[\frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\theta) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta \,|\, \omega_H) - \frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\psi(\theta)) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\psi(\theta) \,|\, \omega_H) \psi'(\theta)\right] \,\mathrm{d}\theta \\ &= \int_{\underline{\theta}}^{\theta^*} \frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\theta) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta \,|\, \omega_H) \,\mathrm{d}\theta + \int_{\theta^*}^{\psi(\underline{\theta})} \frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\theta) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta \,|\, \omega_H) \,\mathrm{d}\theta \\ &\leq \int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial\, \Delta \mathrm{CS}_D}{\partial \theta}(\theta) \frac{\partial F_{\theta\,|\,\omega}}{\partial \omega}(\theta \,|\, \omega_H) \,\mathrm{d}\theta = -\frac{\partial}{\partial \omega} \mathbf{E}\left[\Delta \mathrm{CS}_D(\theta) \,|\, \omega\right]\Big|_{\omega=\omega_H}. \end{aligned}$$

It follows that $\omega \mapsto \mathbf{E}\left[\Delta \mathrm{CS}_D(\theta) \,|\, \omega\right]$ is quasiconcave, as claimed.

(ii) Since $\Delta \mathrm{CS}_I(\theta)$ is increasing in θ if $Q > Q_0$ and $\theta \,|\, \omega_L \succeq_{\mathrm{FOSD}} \theta \,|\, \omega_H$, it follows that $\omega \mapsto \mathbf{E}\left[\Delta \mathrm{CS}_D(\theta) \,|\, \omega\right]$ is non-decreasing in ω if $Q > Q_0$. Similarly, $\omega \mapsto \mathbf{E}\left[\Delta \mathrm{CS}_D(\theta) \,|\, \omega\right]$ is non-increasing in ω if $Q < Q_0$.

A.3 Proofs from Section 4

A.3.1 Proof of Lemma 2

I begin by showing that an allocation function X induces an aggregate quality Q in the market if and only if

$$Q = \int_{\underline{\theta}}^{\overline{\theta}} \left[\delta X(\theta) + q_0(\theta, Q)\left[1 - X(\theta)\right]\right] \, \mathrm{d}F(\theta).$$

The necessity of this condition is obvious. To show that this condition is sufficient, observe that the left-hand side is increasing in Q while the right-hand side is non-increasing in Q. On one hand,

$$\lim_{Q \to 0} \int_{\underline{\theta}}^{\overline{\theta}} \left[\delta X(\theta) + q_0(\theta, Q)\left[1 - X(\theta)\right]\right] \, \mathrm{d}F(\theta) \geq \min\left\{\delta, \lim_{Q \to 0} q_0(\theta, Q)\right\} > 0.$$

On the other hand,

$$\lim_{Q \to +\infty} \int_{\underline{\theta}}^{\overline{\theta}} \left[\delta X(\theta) + q_0(\theta, Q)\left[1 - X(\theta)\right]\right] \, \mathrm{d}F(\theta) \leq \delta + \lim_{Q \to +\infty} \mathbf{E}[q_0(\theta, Q)] < +\infty.$$

Therefore there must be a unique aggregate quality level Q satisfying the condition above; hence the condition is sufficient.

A.3.2 Proof of Lemma 3

Suppose that $x : [\underline{\eta}, \overline{\eta}] \to [0, 1]$ is non-decreasing such that $\operatorname{im} x \subseteq \{0, 1\}$, and that $x = \alpha x_1 + (1 - \alpha)\, x_2$ for $x_1, x_2 \in K$ and $\alpha \in (0, 1)$. Then $\alpha x_1(\eta) + (1 - \alpha)\, x_2(\eta) = x(\eta) \in \{0, 1\}$ for almost every $\eta \in [\underline{\eta}, \overline{\eta}]$, which implies that $x_1(\eta) = x_2(\eta) = x(\eta)$ for almost every $\eta \in [\underline{\eta}, \overline{\eta}]$. Thus $x_1 = x_2$; hence $x \in \operatorname{ex} K$. Conversely, let $x : [\underline{\eta}, \overline{\eta}] \to [0, 1]$ be non-decreasing and suppose the set $\Gamma = \{\eta \in [\underline{\eta}, \overline{\eta}] : x(\eta) \notin \{0, 1\}\}$ has positive measure. Define $x_1, x_2 : [\underline{\eta}, \overline{\eta}] \to [0, 1]$ by $x_1 = x^2$ and $x_2 = 2x - x^2$; by construction, x_1 and x_2 are non-decreasing and $x = (x_1 + x_2)/2$. Note that $x_1 \neq x_2$ (since they are not equal on Γ, which has positive measure). Therefore $x = (x_1 + x_2)/2$ where $x_1, x_2 \in K$ are distinct; hence $x \notin \operatorname{ex} K$.

A.3.3 Completion of the proof of Theorem 1

To apply the results of Bauer (1958) and Szapiel (1975), it remains to show that K is convex and compact. Moreover, the proof of Theorem 1 provided in Section 4.2 assumes the existence of an optimal mechanism; hence it also remains to show that an optimal mechanism exists.

Convexity of K. Let $x_1, x_2 \in K$ for non-decreasing functions $x_1, x_2 : [\underline{\eta}, \overline{\eta}] \to [0,1]$. For any $\alpha \in [0,1]$, observe that $\alpha x_1 + (1-\alpha)\, x_2 : [\underline{\eta}, \overline{\eta}] \to [0,1]$ is increasing. Thus K is convex.

Compactness of K. Since L^1 is a metric space, it suffices to show that K is sequentially compact. Let $\{x_n\}_{n=1}^{\infty} \subseteq K$ be any sequence of functions in K. By the Helly selection theorem, there is a subsequence of $\{x_n\}_{n=1}^{\infty}$ that converges pointwise to some x, which must be a function in K. By the dominated convergence theorem, this subsequence must also converge to x in the the L^1 sense; hence K is compact.

Existence of an optimal mechanism. Say that an effective mechanism (x,t) is *feasible* if it satisfies (IC), (IR), and the equilibrium condition (E). Let $\mathcal{Q} \subseteq \mathbb{R}_+$ be the set of aggregate quality levels in the market that can be induced by a feasible effective mechanism. Observe that the laissez-faire aggregate quality level Q_0 obtains in the market if the policymaker does not intervene (which is always feasible), so $\mathcal{Q}$ is nonempty.

For any sequence of aggregate quality levels $\{Q_n\}_{n=1}^{\infty} \subseteq \mathcal{Q}$ converging to Q, let $\{(x_n, t_n)\}_{n=1}^{\infty}$ be a sequence of feasible effective mechanisms that are optimal for the aggregate quality level Q_n (note that a feasible optimal mechanism conditional on any aggregate quality level $Q_n \in \mathcal{Q}$ always exists). Since uniform transfers between the policymaker and all consumers do not change social welfare, normalize $t_n(\underline{\eta}) = 0$ for all $n = 1, 2, \ldots$. Let (x,t) be the effective mechanism obtained by taking the pointwise limit of $\{(x_n, t_n)\}_{n=1}^{\infty}$; such a limit exists (by passing to a subsequence if necessary) because $0 \le x_n \le 1$ and t_n is bounded as a result of the envelope theorem and the normalization $t_n(\underline{\eta}) = 0$. Observe that p is continuous by Assumption 2; hence Berge's maximum theorem implies $q_0(\theta, \cdot)$ is continuous. It is therefore easy to verify that (x,t) is feasible and induces the aggregate quality level Q in the market.

Let $\Omega^*(Q)$ denote the optimal value of the design problem conditional on the aggregate quality $Q \in \mathcal{Q}$ in the market. By continuity of the policymaker's objective, the above argument shows that $\limsup_{Q_n \to Q} \Omega^*(Q_n) \le \Omega^*(Q)$. Therefore, the supremum of Ω^* must be attained on $\mathcal{Q}$; hence an optimal mechanism exists.

A.3.4 Proof of Proposition 6

While this was proven in Section 4.2, I include a proof here for the sake of completion. When input supply is perfectly elastic, the policymaker's problem can be written as follows:

$$\max_{x} \int_{\underline{\eta}}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta \cdot p(Q_0) + \frac{\int_{\eta}^{\overline{\eta}} \left[\mathbf{E}[\omega \mid s] - 1\right] \, \mathrm{d}G(s)}{g(\eta)}\right] x(\eta) \, \mathrm{d}G(\eta)$$

$$\text{s.t. } x : [\underline{\eta}, \overline{\eta}] \to [0,1] \text{ is non-decreasing.}$$

This is identical to (P) except that the equilibrium constraint (E) no longer has to be included: the policymaker no longer can affect aggregate quality Q through her choice of effective allocation function x. Observe that the objective function is linear and continuous in x, and that K is convex and compact (as shown in the proof of Lemma 2 below). The Bauer maximum principle implies that there exists a maximizer $x^* \in \operatorname{ex} K$. By Lemma 3, each extreme point $x_0 \in \operatorname{ex} K$ satisfies $\operatorname{im} x_0 \subseteq \{0,1\}$. It follows that $\operatorname{im} X^* = \operatorname{im} x^* \subseteq \{0,1\}$, as claimed.

A.3.5 Proof of Proposition 7

Suppose for the sake of contradiction that there exists an optimal mechanism (X^*, T^*) that does not require rationing. Then the corresponding effective allocation function x^*, together with some Lagrange multiplier λ^* and aggregate quality level Q^*, must also solve the auxiliary problem ($\mathrm{P}_{\varnothing}$). Let the corresponding quantity of the public option allocated be denoted by

$$m^* = \int_{\underline{\theta}}^{\overline{\theta}} X^*(\theta) \, \mathrm{d}F(\theta).$$

Let $\eta = u(\delta, \theta) - v_0(\theta, Q^*)$ denote the effective consumption type. By strong duality, the effective allocation function must maximize the policymaker's Lagrangian function:

$$x^* \in \operatorname*{arg\,max}_{x \in K} \int_{\underline{\eta}}^{\overline{\eta}} \lambda^* \left[\mathbf{E}[q_0(\theta, Q^*) \mid \eta] - \delta\right] x(\eta) \, \mathrm{d}G(\eta)$$
$$+ \int_{\underline{\eta}}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta p(Q^*) + \frac{\int_{\eta}^{\overline{\eta}} \left[\mathbf{E}[\omega \mid s] - 1\right] \, \mathrm{d}G(s)}{g(\eta)}\right] x(\eta) \, \mathrm{d}G(\eta).$$

In particular, it must also maximize the Lagrangian function subject to the additional constraint that

$$\int_{\underline{\eta}}^{\overline{\eta}} x(\eta)\ \mathrm{d}G(\eta) = m^*.$$

Since G has full support on $[\underline{\eta}, \overline{\eta}]$, any non-decreasing, right-continuous $x : [\underline{\eta}, \overline{\eta}] \to [0, 1]$ can be represented in quantile space via

$$x(\eta) = \int_0^1 \mathbf{1}_{\eta \geq G^{-1}(1-s)}\ \mathrm{d}\mu(s), \quad \text{for some } \mu \in \Delta([0, 1]).$$

By Fubini's theorem, the policymaker's additionally constrained problem can be rewritten as

$$\max_{\mu \in \Delta([0,1])} \int_0^1 \Psi(s)\ \mathrm{d}\mu(s)$$
$$\text{s.t.} \int_0^1 s\ \mathrm{d}\mu(s) = m^*.$$

Here, the function Ψ is defined by

$$\Psi(s)$$
$$:= \int_{G^{-1}(1-s)}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta p(Q^*) + \frac{\int_\eta^{\overline{\eta}} \left[\mathbf{E}[\omega \mid r] - 1 \right]\ \mathrm{d}G(r)}{g(\eta)} + \lambda^* \left[\mathbf{E}[q_0(\theta, Q^*) \mid \eta] - \delta \right] \right]\ \mathrm{d}G(\eta).$$

Following the results of Aumann and Maschler (1995) and Kamenica and Gentzkow (2011), the value of the policymaker's constrained problem is given by $\operatorname{co}\Psi(m^*)$, where $\operatorname{co}\Psi$ denotes the concave closure of Ψ (i.e., the pointwise smallest concave function that bounds Ψ from above). However, under the assumptions of Proposition 7, $\Psi''(m^*) > 0$, which implies that

$$\Psi(m^*) \neq \operatorname{co}\Psi(m^*).$$

This contradicts the assumption that (X^*, T^*) is an optimal mechanism. It follows that any optimal mechanism (X^*, T^*) requires rationing, as claimed.

A.3.6 Proof of Proposition 8

Given an aggregate quality Q^*, the policymaker's problem can be written as follows:

$$\max_x \int_{\underline{\eta}}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta \cdot p(Q^*) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \mid s] - 1] \ \mathrm{d}G(s)}{g(\eta)} \right] x(\eta) \ \mathrm{d}G(\eta)$$

$$\text{s.t.} \begin{cases} x : [\underline{\eta}, \overline{\eta}] \to [0,1] \text{ is non-decreasing,} \\ \mathbf{E}[q_0(\theta, Q^*)] - Q^* = \displaystyle\int_{\underline{\eta}}^{\overline{\eta}} [\mathbf{E}[q_0(\theta, Q^*) \mid \eta] - \delta] \, x(\eta) \ \mathrm{d}G(\eta). \end{cases}$$

Define

$$H(\eta) := \frac{\int_{\underline{\eta}}^{\eta} [\mathbf{E}[q_0(\theta, Q^*) \mid \eta = s] - \delta] \ \mathrm{d}G(s)}{\mathbf{E}[q_0(\theta, Q^*)] - \delta}.$$

Since the public option induces downward substitution, hence the integrand is always positive for $\underline{\eta} \leq \eta < \overline{\eta}$. Since G has full support on $[\underline{\eta}, \overline{\eta}]$, it follows that H is increasing on $[\underline{\eta}, \overline{\eta}]$. Consequently, one can represent any non-decreasing, right-continuous $x : [\underline{\eta}, \overline{\eta}] \to [0,1]$ in quantile space via

$$x(\eta) = \int_0^1 \mathbf{1}_{\eta \geq H^{-1}(1-s)} \ \mathrm{d}\mu(s), \quad \text{for some } \mu \in \Delta([0,1]).$$

By Fubini's theorem, the policymaker's problem can be rewritten as

$$\max_{\mu \in \Delta([0,1])} \int_0^1 \left[\int_{H^{-1}(1-s)}^{\overline{\eta}} \frac{\eta - c(\delta) - \kappa - \delta \cdot p(Q^*) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \mid s] - 1] \ \mathrm{d}G(s)}{g(\eta)}}{\mathbf{E}[q_0(\theta, Q^*) \mid \eta] - \delta} \ \mathrm{d}H(\eta) \right] \ \mathrm{d}\mu(s)$$

$$\text{s.t.} \int_0^1 s \ \mathrm{d}\mu(s) = \frac{\mathbf{E}[q_0(\theta, Q^*)] - Q^*}{\mathbf{E}[q_0(\theta, Q^*)] - \delta}.$$

Following the results of Aumann and Maschler (1995) and Kamenica and Gentzkow (2011), the value of this problem is given by

$$\text{co}\,\Psi \left(\frac{\mathbf{E}[q_0(\theta, Q^*)] - Q^*}{\mathbf{E}[q_0(\theta, Q^*)] - \delta} \right),$$

where

$$\Psi(s) := \int_{H^{-1}(1-s)}^{\overline{\eta}} \frac{\eta - c(\delta) - \kappa - \delta \cdot p(Q^*) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \mid s] - 1] \ \mathrm{d}G(s)}{g(\eta)}}{\mathbf{E}[q_0(\theta, Q^*) \mid \eta] - \delta} \ \mathrm{d}H(\eta).$$

In particular, the optimal mechanism requires rationing (i.e., the optimal μ^* consists more than a single mass point) if and only if

$$\Psi\left(\frac{\mathbf{E}[q_0(\theta, Q^*)] - Q^*}{\mathbf{E}[q_0(\theta, Q^*)] - \delta}\right) \neq \operatorname{co}\Psi\left(\frac{\mathbf{E}[q_0(\theta, Q^*)] - Q^*}{\mathbf{E}[q_0(\theta, Q^*)] - \delta}\right).$$

A.3.7 Proof of Corollary 1

To obtain a sufficient condition for rationing to be required in the optimal mechanism, observe that a necessary condition for $\Psi(r) = \operatorname{co}\Psi(r)$ is that $\Psi''(r) \leq 0$. As such, rationing is required if $\Psi''(r) > 0$ at $r = (\mathbf{E}[q_0(\theta, Q^*)] - Q^*) \,/\, (\mathbf{E}[q_0(\theta, Q^*)] - \delta)$, where $Q^* \in (\delta, Q_0)$. This condition can be written as

$$\frac{\mathrm{d}}{\mathrm{d}\eta}\left[\frac{\eta - c(\delta) - \kappa - \delta \cdot p(Q^*) + \frac{\int_\eta^{\overline{\eta}}[\mathbf{E}[\omega \mid s] - 1]\ \mathrm{d}G(s)}{g(\eta)}}{\mathbf{E}[q_0(\theta, Q^*) \mid \eta] - \delta}\right]\Bigg|_{\eta = H^{-1}\left(\frac{Q^* - \delta}{\mathbf{E}[q_0(\theta, Q^*)] - \delta}\right)} < 0.$$

Consequently, if any solution $(x_\varnothing, \lambda_\varnothing, Q_\varnothing)$ to the auxiliary problem ($P_\varnothing$) satisfies $\delta < Q_\varnothing < Q_0$ as well as this condition, then the optimal mechanism cannot be a solution to the auxiliary problem ($P_\varnothing$). Therefore, any optimal mechanism must consist of an allocation function in $K \setminus K_\varnothing$; hence it requires rationing.

A.3.8 Proof of Proposition 9

Suppose that the optimal mechanism allocates ΔX units of the public option without rationing. The change in aggregate quality ΔQ induced by the mechanism is determined by

$$Q_0 + \Delta Q = \int_{\underline{\theta}}^{\overline{\theta}} \left[\delta X(\theta) + \left[q_0(\theta, Q_0) + \frac{\partial q_0}{\partial Q}(\theta, Q_0) \cdot \Delta Q\right][1 - X(\theta)]\right]\ \mathrm{d}F(\theta) + o(\Delta Q).$$

This can be rearranged to obtain

$$\Delta Q = \frac{\int_{\underline{\theta}}^{\overline{\theta}} [\delta - q_0(\theta, Q_0)]\, X(\theta)\ \mathrm{d}F(\theta)}{1 - \int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial q_0}{\partial Q}(\theta, Q_0)\,[1 - X(\theta)]\ \mathrm{d}F(\theta)} + o(\Delta Q).$$

When X does not involve rationing, observe that

$$\int_{\underline{\theta}}^{\overline{\theta}} [\delta - q_0(\theta, Q_0)]\, X(\theta)\ \mathrm{d}F(\theta) \to 0 \quad \text{as } \Delta X \to 0.$$

This follows because X must be increasing in $\eta = u(\delta, \theta) - v_0(\theta, Q)$ (cf. Lemma 1); moreover, under the assumption that $q_0(\underline{\theta}, Q_0) < \delta < q_0(\overline{\theta}, Q_0)$,

$$\overline{\eta} = \max_{\theta \in [\underline{\theta}, \overline{\theta}]} [u(\delta, \theta) - v_0(\theta, Q)] \implies \mathbf{E}\left[q_0(\theta, Q) \mid u(\delta, \theta) - v_0(\theta, Q) = \overline{\eta}\right] = \delta.$$

Therefore, $\Delta Q = o(\Delta X)$. Let x be the effective allocation function that represents the allocation function X. Then the change in welfare can be expressed as

$$\int_{\underline{\eta}}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta p(Q_0) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \mid s] - 1]\ \mathrm{d}G(s)}{g(\eta)}\right] x(\eta)\ \mathrm{d}G(\eta) + o(\Delta X).$$

However, observe that

$$\overline{\eta} = c(\delta) - \delta p(Q_0) + o(\Delta X).$$

Consequently, it follows that the change in welfare can be rewritten as

$$-\kappa \cdot \Delta X + o(\Delta X).$$

As $\kappa > 0$ by assumption, the change in welfare arising from the public option is negative; in turn, this contradicts the optimality of the mechanism. By continuity, this argument extends to any sufficiently small ΔX. Therefore, if the optimal mechanism allocates a sufficiently small amount of the public option, then it must involve rationing.

A.4 Proofs from Section 5

A.4.1 Proof of Theorem 2

The marginal effect of quality on total weighted surplus can be straightforwardly derived by using the envelope theorem (cf. Lemma 1) and implicit differentiation. This computation is routine and therefore omitted.

A.4.2 Proof of Proposition 10

Under the sufficient condition given in Proposition 10, Theorem 2 implies that the marginal effect of quality on total weighted surplus at any policy (δ, X, T) that induces an aggregate quality level of Q can be written as

$$\frac{\partial \mathrm{W}}{\partial \delta} = \int_{\underline{\theta}}^{\overline{\theta}} \left[\left[\theta - \frac{\int_{\theta}^{\overline{\theta}} [1 - \mathbf{E}[\omega \,|\, s]] \ \mathrm{d}F(s)}{f(\theta)} \right] \nu'(\delta) - c'(\delta) - p(Q) \right] X(\theta) \ \mathrm{d}F(\theta)$$
$$+ \left[\alpha Q - \int_{\underline{\theta}}^{\overline{\theta}} [\delta X(\theta) + \mathbf{E}[\omega \,|\, \theta] q_0(\theta, Q) \, [1 - X(\theta)]] \ \mathrm{d}F(\theta) \right] \cdot \frac{\partial P}{\partial \delta}.$$

Since $\nu'(\delta) \to +\infty$ as $\delta \to 0$ by Assumption 1, the sufficient condition given in Proposition 10 ensures that $\partial \mathrm{W}/\partial \delta \to +\infty$ as $\delta \to 0$ for any incentive-compatible mechanism. In particular, this must also hold for the optimal mechanism at $\delta = 0$. This implies that $\delta^* \neq 0$ because the policymaker could otherwise always strictly increase total weighted surplus by increasing quality.

A.4.3 Proof of Proposition 11

Under the sufficient condition given in Proposition 11, the marginal effect of quality on total weighted surplus at the optimal policy (δ^*, X^*, T^*) that induces an aggregate quality level of Q^* can be written as

$$\frac{\partial \mathrm{W}}{\partial \delta} = \int_{\underline{\theta}}^{\overline{\theta}} \left[\left[\theta - \frac{\int_{\theta}^{\overline{\theta}} [1 - \mathbf{E}[\omega \,|\, s]] \ \mathrm{d}F(s)}{f(\theta)} \right] \nu'(\delta^*) - c'(\delta^*) - p(Q^*) \right] X^*(\theta) \ \mathrm{d}F(\theta)$$
$$- \int_{\underline{\theta}}^{\overline{\theta}} [\delta^* X^*(\theta) + \mathbf{E}[\omega \,|\, \theta] q_0(\theta, Q^*) \, [1 - X^*(\theta)]] \ \mathrm{d}F(\theta) \cdot \frac{\partial P}{\partial \delta}.$$

The sufficient condition given in Proposition 11 ensures that $\partial \mathrm{W}/\partial \delta < 0$. Since (δ^*, X^*, T^*) is the optimal policy, it follows that $\delta^* = 0$ because the policymaker could otherwise always strictly increase total weighted surplus by decreasing quality.

A.5 Proofs from Section 6

A.5.1 Proof of Theorem 3

The existence of an optimal mechanism (X^*, T^*) can be proven in a similar way as in Theorem 1 and omitted here for brevity. Suppose that the optimal mechanism induces an aggregate quality of Q and an aggregate externality of E. Similar to (P), the policymaker's problem can be written as

$$\max_x \int_{\underline{\eta}}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta \cdot p(Q) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \mid s] - 1] \, \mathrm{d}G(s)}{g(\eta)}\right] x(\eta) \, \mathrm{d}G(\eta)$$

$$\text{s.t.} \begin{cases} x : [\underline{\eta}, \overline{\eta}] \to [0,1] \text{ is non-decreasing}, \\ Q - \mathbf{E}[q_0(\theta, Q)] = \displaystyle\int_{\underline{\eta}}^{\overline{\eta}} [\delta - \mathbf{E}[q_0(\theta, Q) \mid \eta]] \, x(\eta) \, \mathrm{d}G(\eta), \\ E = \displaystyle\int_{\underline{\eta}}^{\overline{\eta}} \mathbf{E}[\xi \mid \eta] \, [\delta x(\eta) + \mathbf{E}[q_0(\theta, Q) \mid \eta] \, [1 - x(\eta)]] \, \mathrm{d}G(\eta). \end{cases}$$

Notice that the policymaker's objective function is identical to that in (P); it is linear (hence convex) and continuous in x, so that the results of Bauer (1958) and Szapiel (1975) apply. Define the function $\ell : K \to \mathbb{R}^2$ and the set $\Sigma \subseteq \mathbb{R}^2$ by

$$\begin{cases} \ell(x) & := \left(\displaystyle\int_{\underline{\eta}}^{\overline{\eta}} [\delta - \mathbf{E}[q_0(\theta, Q) \mid \eta]] \, x(\eta) \, \mathrm{d}G(\eta), \int_{\underline{\eta}}^{\overline{\eta}} \mathbf{E}\left[\xi \left[\delta - q_0(\theta, Q)\right] \mid \eta\right] x(\eta) \, \mathrm{d}G(\eta)\right), \\ \Sigma & := \{Q - \mathbf{E}[q_0(\theta, Q)], E - \mathbf{E}[\xi q_0(\theta, Q)]\}. \end{cases}$$

Clearly, ℓ is continuous and linear, and Σ is closed and convex. Thus the optimal effective allocation function x^* can be written as the convex combination of at most three extreme points of K. By Lemma 3, there exist $0 < \pi_1 < \pi_2 < 1$ such that $\operatorname{im} x^* \subseteq \{0, \pi_1, \pi_2, 1\}$.

Finally, if the externality function e is convex, notice that the policymaker's problem can then be written as

$$\max_x \int_{\underline{\eta}}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta \cdot p(Q) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \mid s] - 1] \, \mathrm{d}G(s)}{g(\eta)}\right] x(\eta) \, \mathrm{d}G(\eta)$$
$$+ e\left(\int_{\underline{\eta}}^{\overline{\eta}} \mathbf{E}[\xi \mid \eta] \, [\delta x(\eta) + \mathbf{E}[q_0(\theta, Q) \mid \eta] \, [1 - x(\eta)]] \, \mathrm{d}G(\eta)\right)$$

$$\text{s.t.} \begin{cases} x : [\underline{\eta}, \overline{\eta}] \to [0,1] \text{ is non-decreasing}, \\ Q - \mathbf{E}[q_0(\theta, Q)] = \displaystyle\int_{\underline{\eta}}^{\overline{\eta}} [\delta - \mathbf{E}[q_0(\theta, Q) \mid \eta]] \, x(\eta) \, \mathrm{d}G(\eta). \end{cases}$$

Crucially, observe that the objective function in this problem is nonetheless continuous and convex in x, even if it is no longer linear. Thus the results of Bauer (1958) and Szapiel (1975) still apply;

like Theorem 1, at most two prices are required in this case.

A.5.2 Proof of Theorem 4

The existence of an optimal mechanism (X^*, T^*) can be proven in a similar way as in Theorem 1 and omitted here for brevity. Instead, I begin by showing the implications of Assumption 3.

Lemma 4. *Under Assumptions 1 and 3, there exists a non-increasing function $\phi : \mathbb{R}_+ \to \mathbb{R}_+$ such that consumer demand for quality $D(p, \theta)$ can be written in the form*

$$D(p, \theta) = \theta\phi(p).$$

Note that under regularity conditions on ϕ, a converse to Lemma 4 can actually be shown. However, as that direction is not relevant to the proof of Theorem 4, I do not pursue it here.

Proof. As $u(\cdot; \theta)$ is strictly concave, the solution to each consumer's utility maximization problem is uniquely given by

$$D(p, \theta) \in \arg\max_{q \in \mathbb{R}_+} \left[u(q, \theta) - pq\right].$$

Therefore $D(p, \theta)$ satisfies the first-order condition

$$\frac{\partial u}{\partial q}(D(p, \theta), \theta) = p.$$

Since u is a homogeneous function of degree 1 in (q, θ), its partial derivatives must be homogeneous functions of degree 0 in (q, θ) as a consequence of Euler's theorem. As such, for any $a > 0$ and $\theta \in [\underline{\theta}, \overline{\theta}]$,

$$\frac{\partial u}{\partial q}\left(\frac{D(p, a\theta)}{a}, \theta\right) = \frac{\partial u}{\partial q}(D(p, a\theta), a\theta) = p = \frac{\partial u}{\partial q}(D(p, \theta), \theta).$$

Because u is strictly concave, this implies that

$$D(p, a\theta) = aD(p, \theta) \implies D(p, \theta) = \theta \cdot \frac{D(p, \underline{\theta})}{\underline{\theta}}.$$

Thus the required non-increasing function ϕ can be obtained by setting $\phi(p) = D(p, \underline{\theta})/\underline{\theta}$. □

Next, I combine the result of Lemma 4 and Assumption 3 to obtain a new equilibrium condition (analogous to Lemma 2) when the input market is operated by a monopolist.

Lemma 5. *Suppose that consumer demand for quality can be written in the form* $D(p,\theta) = \theta\phi(p)$*. An allocation function* X *induces a price* p *per unit of quality in the market if and only if the following equilibrium condition is satisfied:*

$$\frac{\int_{\underline{\theta}}^{\overline{\theta}} [\delta - \theta\phi(p)]\, X(\theta)\ \mathrm{d}F(\theta)}{\int_{\underline{\theta}}^{\overline{\theta}} \theta\,[1 - X(\theta)]\ \mathrm{d}F(\theta)} + p + \frac{\phi(p)}{\phi'(p)} = C'\left(\int_{\underline{\theta}}^{\overline{\theta}} [\delta X(\theta) + \theta\phi(p)\,[1 - X(\theta)]]\ \mathrm{d}F(\theta)\right).$$

Proof. Given the allocation function X, the monopolist solves

$$\max_p \left\{ p \int_{\underline{\theta}}^{\overline{\theta}} [\delta X(\theta) + \theta\phi(p)\,[1 - X(\theta)]]\ \mathrm{d}F(\theta) - C\left(\int_{\underline{\theta}}^{\overline{\theta}} [\delta X(\theta) + \theta\phi(p)\,[1 - X(\theta)]]\ \mathrm{d}F(\theta)\right)\right\}.$$

The monopolist's first-order condition can be rewritten as

$$\frac{\int_{\underline{\theta}}^{\overline{\theta}} [\delta - \theta\phi(p)]\, X(\theta)\ \mathrm{d}F(\theta)}{\int_{\underline{\theta}}^{\overline{\theta}} \theta\,[1 - X(\theta)]\ \mathrm{d}F(\theta)} + p + \frac{\phi(p)}{\phi'(p)} = C'\left(\int_{\underline{\theta}}^{\overline{\theta}} [\delta X(\theta) + \theta\phi(p)\,[1 - X(\theta)]]\ \mathrm{d}F(\theta)\right).$$

It remains to show that there exists a unique p that satisfies this condition. To this end, note that the right-hand side is non-increasing in p, while Assumption 3 guarantees that the left-hand side is increasing in p as

$$\begin{aligned} p \mapsto D(p,\underline{\theta}) \quad \text{is log-concave} &\iff p \mapsto \log\phi(p) \quad \text{is concave} \\ &\iff p \mapsto \frac{\phi(p)}{\phi'(p)} \quad \text{is non-decreasing.} \end{aligned}$$

□

Finally, Lemma 5 allows the optimal mechanism to be characterized by conditioning on the optimal price p per unit quality induced by the optimal mechanism. With the change of variables

$\eta = u(\delta, \theta) - v_0(\theta, \mathbf{E}[\theta] \cdot \phi(p))$, the policymaker's problem can be written as

$$\max_x \int_{\underline{\eta}}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta \cdot p + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \mid s] - 1] \, \mathrm{d}G(s)}{g(\eta)} \right] x(\eta) \, \mathrm{d}G(\eta)$$

$$\text{s.t.} \begin{cases} x : [\underline{\eta}, \overline{\eta}] \to [0,1] \text{ is non-decreasing,} \\ \displaystyle\int_{\underline{\eta}}^{\overline{\eta}} \mathbf{E}[\theta \mid \eta] x(\eta) \, \mathrm{d}G(\eta) = \mathbf{E}[\theta] - \frac{Q - \mathbf{E}[\theta] \cdot \phi(p)}{C'(Q) - p - \frac{\phi(p)}{\phi'(p)}}, \\ \displaystyle\int_{\underline{\eta}}^{\overline{\eta}} [\delta - \mathbf{E}[\theta \mid \eta] \phi(p)] \, x(\eta) \, \mathrm{d}G(\eta) = Q - \mathbf{E}[\theta] \cdot \phi(p). \end{cases}$$

Observe that, unlike (E) in Lemma 2, the equilibrium condition in Lemma 5 is not affine in the allocation function X. However, the equilibrium condition consists of two constituent expressions, each of which is affine in the allocation function X—and hence the effective allocation function x. Conditioning on the values of these constituent expressions then allows the optimal mechanism to be characterized.

To this end, notice that the policymaker's objective function is identical to that in (P); it is linear (hence convex) and continuous in x, so that the results of Bauer (1958) and Szapiel (1975) apply. Define the function $\ell : K \to \mathbb{R}^2$ and the set $\Sigma \subseteq \mathbb{R}^2$ by

$$\begin{cases} \ell(x) & := \left(\displaystyle\int_{\underline{\eta}}^{\overline{\eta}} \mathbf{E}[\theta \mid \eta] x(\eta) \, \mathrm{d}G(\eta), \int_{\underline{\eta}}^{\overline{\eta}} [\delta - \mathbf{E}[\theta \mid \eta] \phi(p)] \, x(\eta) \, \mathrm{d}G(\eta) \right), \\ \Sigma & := \left\{ \mathbf{E}[\theta] - \dfrac{Q - \mathbf{E}[\theta] \cdot \phi(p)}{C'(Q) - p - \frac{\phi(p)}{\phi'(p)}}, Q - \mathbf{E}[\theta] \cdot \phi(p) \right\}. \end{cases}$$

Clearly, ℓ is continuous and linear, and Σ is closed and convex. Thus the optimal effective allocation function x^* can be written as the convex combination of at most three extreme points of K. By Lemma 3, there exist $0 < \pi_1 < \pi_2 < 1$ such that $\operatorname{im} x^* \subseteq \{0, \pi_1, \pi_2, 1\}$.

Appendix B Computation of the optimal mechanism

In this appendix, I supplement the non-constructive characterization of the optimal mechanism (cf. Theorem 1) by developing a method to compute the optimal mechanism.

The key idea of this method is to augment the linear program faced by the policymaker's problem into a quadratic program. To this end, recall the policymaker's problem (P) introduced in Section 4.2:

$$\max_x \int_{\underline{\eta}}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta \cdot p(Q) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \mid s] - 1] \ \mathrm{d}G(s)}{g(\eta)}\right] x(\eta) \ \mathrm{d}G(\eta)$$

$$\text{s.t.} \begin{cases} x : [\underline{\eta}, \overline{\eta}] \to [0,1] \text{ is non-decreasing,} \\ Q - \mathbf{E}[q_0(\theta, Q)] = \displaystyle\int_{\underline{\eta}}^{\overline{\eta}} [\delta - \mathbf{E}[q_0(\theta, Q) \mid \eta]]\, x(\eta) \ \mathrm{d}G(\eta). \end{cases}$$

Augmentation is done by including a term that depends on the square of the effective allocation function. Specifically, the augmented problem is

$$\max_x \int_{\underline{\eta}}^{\overline{\eta}} \left[\left[\eta - c(\delta) - \kappa - \delta \cdot p(Q) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \mid s] - 1] \ \mathrm{d}G(s)}{g(\eta)}\right] x(\eta) - \frac{\varepsilon}{2} [x(\eta)]^2\right] \ \mathrm{d}G(\eta)$$

$$\text{s.t.} \begin{cases} x : [\underline{\eta}, \overline{\eta}] \to [0,1] \text{ is non-decreasing,} \\ Q - \mathbf{E}[q_0(\theta, Q)] = \displaystyle\int_{\underline{\eta}}^{\overline{\eta}} [\delta - \mathbf{E}[q_0(\theta, Q) \mid \eta]]\, x(\eta) \ \mathrm{d}G(\eta). \end{cases}$$

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

Letting $\lambda \in \mathbb{R}$ denote the Lagrange multiplier for the equilibrium constraint, the Lagrangian function for the augmented problem is

$$\begin{aligned} L_\varepsilon(x; Q, \lambda) = & \int_{\underline{\eta}}^{\overline{\eta}} \left[\left[\eta - c(\delta) - \kappa - \delta \cdot p(Q) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \mid s] - 1] \ \mathrm{d}G(s)}{g(\eta)}\right] x(\eta) - \frac{\varepsilon}{2} [x(\eta)]^2\right] \ \mathrm{d}G(\eta) \\ & + \lambda \left[Q - \mathbf{E}[q_0(\theta, Q)] - \int_{\underline{\eta}}^{\overline{\eta}} [\delta - \mathbf{E}[q_0(\theta, Q) \mid \eta]]\, x(\eta) \ \mathrm{d}G(\eta)\right]. \end{aligned}$$

For ease of notation, define

$$\phi_\varepsilon(\eta; Q, \lambda) := \frac{1}{\varepsilon}\left[\eta - c(\delta) - \kappa - \delta p(Q) - \lambda\left[\delta - \mathbf{E}[q_0(\theta, Q) \mid \eta]\right] + \frac{\int_\eta^{\overline{\eta}} \left[\mathbf{E}[\omega \mid s] - 1\right] \, \mathrm{d}G(s)}{g(\eta)}\right].$$

Lemma 6. *Fix $\varepsilon > 0$ and the Lagrange multiplier $\lambda_\varepsilon \in \mathbb{R}$. Then there exists a unique solution to the maximization problem*

$$\max_x \int_{\underline{\eta}}^{\overline{\eta}} \varepsilon \left[\phi_\varepsilon(\eta; Q, \lambda_\varepsilon) \cdot x(\eta) - \frac{1}{2}\left[x(\eta)\right]^2\right] \, \mathrm{d}G(\eta)$$

$$\text{s.t. } x : [\underline{\eta}, \overline{\eta}] \to [0,1] \text{ is non-decreasing.}$$

Proof. Define

$$K := \left\{x : [\underline{\eta}, \overline{\eta}] \to [0,1] \text{ is non-decreasing}\right\} \subseteq L^2([\underline{\eta}, \overline{\eta}]).$$

Observe that K is nonempty, compact, and convex (as in the proof of Theorem 1). Moreover, by completing the square, the above problem can be rewritten as (by omitting additive terms that do not depend on x)

$$\frac{\varepsilon}{2} \cdot \min_{x \in K} \int_{\underline{\eta}}^{\overline{\eta}} \left[x(\eta) - \phi_\varepsilon(\eta; Q, \lambda_\varepsilon)\right]^2 \, \mathrm{d}G(\eta).$$

By the Hilbert projection theorem, a unique solution x^* exists to the minimization problem that is equal to the projection of $\phi_\varepsilon(\cdot; Q, \lambda_\varepsilon)$ onto K. □

Next, I show that the projection of $\phi_\varepsilon(\cdot; Q, \lambda_\varepsilon)$ onto K in the proof of Lemma 6 takes a simple form for any $\varepsilon > 0$. This is given by Lemma 2 of my companion paper (Kang, 2022), which I include here for the sake of completeness.

Lemma 7. *Let $\mathcal{Y} := \{y : [0,1] \to [0,1]$ is non-decreasing$\}$ and let $\chi : [0,1] \to \mathbb{R}$ be a square-integrable function. Then the unique solution to the problem*

$$\min_{y \in \mathcal{Y}} \int_0^1 \left[y(s) - \chi(s)\right]^2 \, \mathrm{d}s$$

is

$$y^*(s) := \begin{cases} 0 & \text{if } \overline{\chi}(s) \leq 0, \\ 1 & \text{if } \overline{\chi}(s) \geq 1, \\ \overline{\chi}(s) & \text{otherwise,} \end{cases} \quad \text{where } \overline{\chi}(s) := -\frac{\mathrm{d}}{\mathrm{d}s}\left(\mathrm{co} \int_s^1 \chi(r) \, \mathrm{d}r\right).$$

Proof of Lemma 7. First, assume that $\overline{\chi}(s) > 0$ for all $s \in [0,1]$. Define

$$\mathcal{I} := \{y : [0,1] \to \mathbb{R} \text{ is non-decreasing}\}.$$

Let $\Pi_{\mathcal{I}}\chi$ denote the projection of χ onto $\mathcal{I}$. Since $\mathcal{I} \subseteq L^2([0,1])$ is also nonempty, closed, and convex, the projection is unique, hence $\Pi_{\mathcal{I}}\chi$ is well-defined. Since $\mathcal{I}$ is a convex cone,

$$\int_0^1 [\Pi_{\mathcal{I}}\chi(s) - \chi(s)]\, y(s)\ \mathrm{d}s \geq 0 \quad \text{for any } y \in \mathcal{I}.$$

In particular, choose $y(s) = \mathbf{1}_{s>r}$ for some $r \in [0,1]$. Then the above implies that

$$\overline{Y}(r) := \int_r^1 \Pi_{\mathcal{I}}\chi(s)\ \mathrm{d}s \geq \int_r^1 \chi(s)\ \mathrm{d}s =: Y(r).$$

Since $\Pi_{\mathcal{I}}\chi$ is non-decreasing on $[0,1]$, $\overline{Y}$ must be concave; hence $\overline{Y}$ is a concave majorant of Y. Now, if $\overline{Y}$ is not the *least* concave majorant of Y, then there exist $0 < s_1 < s_2 < 1$ and a separating linear function $\ell(s)$, such that $\ell(s) \geq Y(s)$ for $s \in [0,1]$; $\overline{Y}(s) > \ell(s)$ for $s \in (s_1, s_2)$; and $\overline{Y}(s_i) = \ell(s_i)$ for $i = 1, 2$. Define $\widetilde{\chi}$ by $\widetilde{\chi}(s) = \Pi_{\mathcal{I}}\chi(s)$ for $s \notin (s_1, s_2)$ and $\widetilde{\chi}(s) = \ell'(s)$ for $s \in (s_1, s_2)$. Then integration by parts yields the following contradiction:

$$\begin{aligned} 0 &\leq \int_0^1 [\Pi_{\mathcal{I}}\chi(s) - \chi(s)]\, [\widetilde{\chi}(s) - \Pi_{\mathcal{I}}\chi(s)]\ \mathrm{d}s \\ &= -\int_0^1 \left[Y(s) - \overline{Y}(s)\right]\ \mathrm{d}\left[\widetilde{\chi}(s) - \Pi_{\mathcal{I}}\chi(s)\right] = \int_{s_1}^{s_2} \left[Y(s) - \overline{Y}(s)\right] \mathrm{d}\Pi_{\mathcal{I}}\chi(s) < 0. \end{aligned}$$

Therefore $\overline{Y}$ is the least concave majorant of Y. It follows that $\Pi_{\mathcal{I}}\chi = -(\operatorname{co} Y)' = \overline{\chi}$.

It remains to consider the case where either $\overline{\chi}(s) \leq 0$ or $\overline{\chi}(s) \geq 1$ for some $s \in [0,1]$. Define ρ so that $\overline{\chi}(s) \leq 0$ for all $s \in [0, \rho)$; and $0 < \overline{\chi}(s)$ for all $s \in (\rho, 1]$. Such a $\rho \in [0,1]$ exists since $\overline{\chi}$ is non-decreasing by construction. The argument above shows that $\overline{\chi}$ minimizes the integral between ρ and 1:

$$\overline{\chi} \in \arg\min_{y \in \mathcal{Y}} \int_\rho^1 [y(s) - \chi(s)]^2\ \mathrm{d}s.$$

Since $\mathcal{Y} \subset \mathcal{I}$, the argument above also shows that the constraint $y(s) \geq 0$ must bind for $s \in [0, \rho]$. The case where $\overline{\chi}(s) \geq 1$ can be similarly dealt with. This yields the solution y^* as claimed. □

I now use Lemma 7 to solve the augmented problem. To do so, I employ the following change of variables:

$$y(s) := x(G^{-1}(s)) \quad \text{and} \quad \chi(s) := \phi_\varepsilon(G^{-1}(s); Q, \lambda_\varepsilon) \quad \text{for every } s \in [0,1].$$

Then, using the result of Lemma 7, define

$$\begin{aligned} \overline{\chi}(s) &:= -\frac{\mathrm{d}}{\mathrm{d}s}\left(\operatorname{co}\int_s^1 \chi(r)\ \mathrm{d}r\right) \\ &= -\frac{\mathrm{d}}{\mathrm{d}s}\left(\operatorname{co}\int_s^1 \phi_\varepsilon(G^{-1}(r); Q, \lambda_\varepsilon)\ \mathrm{d}r\right) \quad \text{for every } s \in [0,1]. \end{aligned}$$

Consequently,

$$\overline{\chi}(G(\eta)) = \frac{\mathrm{d}}{\mathrm{d}s}\left(\operatorname{co}\int_{1-s}^1 \phi_\varepsilon(G^{-1}(r); Q, \lambda_\varepsilon)\ \mathrm{d}r\right)\bigg|_{s=1-G(\eta)} \qquad \text{for every } \eta \in [\underline{\eta}, \overline{\eta}].$$

To reverse the change of variables, define $\overline{\phi}_\varepsilon(\cdot; Q, \lambda_\varepsilon) := \overline{\chi} \circ G$ and observe that

$$\begin{aligned} \min_{y \in \mathcal{Y}} \int_0^1 [y(t) - \chi(s)]^2\ \mathrm{d}s &= \min_{x \in K} \int_0^1 \left[x(G^{-1}(s)) - \phi_\varepsilon(G^{-1}(s); Q, \lambda_\varepsilon)\right]^2\ \mathrm{d}s \\ &= \min_{x \in K} \int_{\underline{\eta}}^{\overline{\eta}} [x(\eta) - \phi_\varepsilon(\eta; Q, \lambda_\varepsilon)]^2\ \mathrm{d}G(\eta). \end{aligned}$$

Therefore, Lemma 7 implies that the unique optimal allocation function for the augmented problem is given by

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

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

Theorem 5. *As $\varepsilon \searrow 0$, x_ε^* converges pointwise to an optimal effective allocation function x^*.*

Proof. Because $\{x_{1/n}^*\}_{n=1}^\infty$ is uniformly bounded, Helly's selection theorem applies: there exists a subsequence of $\{x_{1/n}^*\}_{n=1}^\infty$ that converges pointwise to some function x^*. Since convergence is pointwise and each $x_{1/n}^*$ in the subsequence is feasible, hence x^* must also be feasible.

Suppose on the contrary that x^* is not a solution of the policymaker's problem (P). Then there exists an effective allocation function x° that satisfies the equilibrium condition (E), such that

$$\int_{\underline{\eta}}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta \cdot p(Q) + \frac{\int_{\eta}^{\overline{\eta}} \left[\mathbf{E}[\omega \mid s] - 1 \right] \ \mathrm{d}G(s)}{g(\eta)} \right] x^\circ(\eta) \ \mathrm{d}G(\eta)$$
$$> \int_{\underline{\eta}}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta \cdot p(Q) + \frac{\int_{\eta}^{\overline{\eta}} \left[\mathbf{E}[\omega \mid s] - 1 \right] \ \mathrm{d}G(s)}{g(\eta)} \right] x^*(\eta) \ \mathrm{d}G(\eta).$$

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

$$\int_{\underline{\eta}}^{\overline{\eta}} \left[\left[\eta - c(\delta) - \kappa - \delta \cdot p(Q) + \frac{\int_{\eta}^{\overline{\eta}} \left[\mathbf{E}[\omega \mid s] - 1 \right] \ \mathrm{d}G(s)}{g(\eta)} \right] x^\circ(\eta) - \frac{\varepsilon}{2} \left[x^\circ(\eta) \right]^2 \right] \ \mathrm{d}G(\eta)$$
$$> \int_{\underline{\eta}}^{\overline{\eta}} \left[\left[\eta - c(\delta) - \kappa - \delta \cdot p(Q) + \frac{\int_{\eta}^{\overline{\eta}} \left[\mathbf{E}[\omega \mid s] - 1 \right] \ \mathrm{d}G(s)}{g(\eta)} \right] x^*(\eta) - \frac{\varepsilon}{2} \left[x^*(\eta) \right]^2 \right] \ \mathrm{d}G(\eta).$$

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

It remains to show that the entire sequence of $\{x_{1/n}^*\}_{n=1}^\infty$ (i.e., not just a subsequence) converges to x^*. This follows from the continuity of x_ε^* in ε due to the continuity of the projection operator.

□

Appendix C Additional results and discussion

In this appendix, I show how the trade-off between the direct and indirect effects determines when to use a public option in the first place.

I begin by demonstrating how the optimality of a public option can depend on local properties of input supply.

Proposition 12. *Let Q_0 denote the laissez-faire aggregate quality level, and suppose that*

$$\mathbf{E}\left[\omega\left[u(\delta,\theta)-v_0(\theta,Q_0)\right]\right]+\frac{\left[\mathbf{E}[\omega q_0(\theta,Q_0)]-\alpha Q_0\right]\left[\mathbf{E}[q_0(\theta,Q_0)]-\delta\right]p'(Q_0)}{1-\mathbf{E}\left[\frac{\partial q_0}{\partial Q}(\theta,Q_0)\right]}>\kappa+c(\delta)+\delta p(Q_0).$$

Then a public option is optimal; that is, any optimal allocation function X^* satisfies $X^*\not\equiv 0$.

Proof. Consider the family of allocation functions $X(\theta)\equiv\pi$, where $\pi\in[0,1]$. Then the equilibrium condition (E) can be written as

$$Q-\mathbf{E}[q_0(\theta,Q)]=\left[\delta-\mathbf{E}[q_0(\theta,Q)]\right]\pi.$$

The increase in aggregate quality $\mathrm{d}Q$ that arises from an infinitesimal increase $\mathrm{d}\pi$ in allocation probability for all consumers is given by

$$\left[1-(1-\pi)\,\mathbf{E}\left[\frac{\partial q_0}{\partial Q}(\theta,Q)\right]\right]\cdot\mathrm{d}Q=\left[\delta-\mathbf{E}[q_0(\theta,Q)]\right]\cdot\mathrm{d}\pi.$$

Recall that total social welfare is given by

$$\left[\mathbf{E}\left[\omega\left[u(\delta,\theta)-v_0(\theta,Q)\right]\right]-c(\delta)-\kappa-\delta p(Q)\right]\pi+\alpha\,\mathrm{PS}(Q)+\mathbf{E}\left[\omega v_0(\theta,Q)\right].$$

Thus the change in total social welfare that arises from an infinitesimal increase $\mathrm{d}\pi$ in allocation probability, relative to $\pi=0$ where $Q=Q_0$, is

$$\begin{aligned}&\left[\mathbf{E}\left[\omega\left[u(\delta,\theta)-v_0(\theta,Q_0)\right]\right]-c(\delta)-\kappa-\delta p(Q_0)\right]\cdot\mathrm{d}\pi\\&\quad+\frac{\left[\mathbf{E}\left[\omega q_0(\theta,Q_0)\right]-\alpha Q_0\right]\left[\mathbf{E}\left[q_0(\theta,Q_0)\right]-\delta\right]p'(Q_0)}{1-\mathbf{E}\left[\frac{\partial q_0}{\partial Q}(\theta,Q_0)\right]}\cdot\mathrm{d}\pi.\end{aligned}$$

This is strictly positive when the condition in Proposition 12 holds:

$$\mathbf{E}\left[\omega\left[u(\delta,\theta)-v_0(\theta,Q_0)\right]\right]+\frac{\left[\mathbf{E}[\omega q_0(\theta,Q_0)]-\alpha Q_0\right]\left[\mathbf{E}[q_0(\theta,Q_0)]-\delta\right]p'(Q_0)}{1-\mathbf{E}\left[\frac{\partial q_0}{\partial Q}(\theta,Q_0)\right]}>\kappa+c(\delta)+\delta p(Q_0).$$

This concludes the proof of Proposition 12. □

Using an envelope theorem argument, Proposition 12 establishes a sufficient condition for a public option to be optimal based on the trade-off between its direct and indirect effects. On one hand, when input supply is perfectly elastic, then the condition in Proposition 12 reduces to

$$\mathbf{E}\left[\omega\left[u(\delta,\theta)-v_0(\theta,Q_0)\right]\right]>\kappa+c(\delta)+\delta p(Q_0).$$

But this condition is never satisfied as long as the distribution F is not trivial (i.e., when $\underline{\theta}\neq\overline{\theta}$). (Note, however, that this is a sufficient condition: a public option might still be optimal even when $p'(Q_0)=0$.) On the other hand, when input supply is not perfectly elastic, then the condition in Proposition 12 is satisfied when the indirect effect is sufficiently large.

The sufficient condition given in Proposition 12 is most easily understood through a parametric example. Suppose that consumer utility is given by $u(q,\theta)=\theta^{1/\varepsilon}q^{1-1/\varepsilon}$, so that $\varepsilon>1$ denotes each consumer's (constant) elasticity of demand for quality. Suppose that there is costless conversion from input into the final good so that $c\equiv 0$, and let input supply be $p(Q)=AQ^{1/\xi}$, where

$$A=\frac{\varepsilon-1}{\varepsilon}\left(\frac{Q_0}{\mathbf{E}[\theta]}\right)^{-1/\varepsilon}Q_0^{-1/\xi}.$$

This parametrization ensures that the laissez-faire aggregate quality level as ε and ξ vary remains at Q_0, where ξ is the (constant) elasticity of input supply. Then the condition can be written as

$$\mathbf{E}\left[\omega\left[\theta^{1/\varepsilon}\delta^{1-1/\varepsilon}-\frac{\theta}{\varepsilon}\left(\frac{Q_0}{\mathbf{E}[\theta]}\right)^{1-1/\varepsilon}\right]\right]+\left(1-\frac{1}{\varepsilon}\right)\left(\frac{Q_0}{\mathbf{E}[\theta]}\right)^{-1/\varepsilon}\left[\frac{(Q_0-\delta)\left(\frac{\mathbf{E}[\omega\theta]}{\mathbf{E}[\theta]}-\alpha\right)}{\varepsilon+\xi}-\delta\right]>\kappa.$$

While the magnitude of the indirect effect need not be monotone in ε, it is decreasing in ξ; hence it is larger when input supply is more inelastic. When α is sufficiently small and $\delta<Q_0$, then a more inelastic supply also implies that the condition is more likely to be satisfied.

I conclude by showing how the optimality of a public option can depend on global properties of input supply.

Proposition 13. *Let Q_0 denote the laissez-faire aggregate quality level, and suppose that there exists Q such that $\delta < Q < Q_0$ and*

$$\mathbf{E}\left[\omega\left[v_0(\theta, Q) - v_0(\theta, Q_0)\right]\right] + \mathbf{E}\left[\omega\left[u(\delta, \theta) - v_0(\theta, Q)\right]\right] \cdot \frac{\mathbf{E}[q_0(\theta, Q)] - Q}{\mathbf{E}[q_0(\theta, Q)] - \delta}$$
$$> \kappa + [c(\delta) + \delta p(Q)] \cdot \frac{\mathbf{E}[q_0(\theta, Q)] - Q}{\mathbf{E}[q_0(\theta, Q)] - \delta} + \alpha\left[\mathrm{PS}(Q_0) - \mathrm{PS}(Q)\right].$$

Then a public option is optimal; that is, any optimal allocation function X^* satisfies $X^* \not\equiv 0$.

Proof. For any given aggregate quality level Q, denote the value of the policymaker's problem (P) by $\mathrm{val}(Q)$. A public option is optimal if

$$\mathrm{val}(Q) + \mathbf{E}\left[\omega\left[v_0(\theta, Q) - v_0(\theta, Q_0)\right]\right] + \alpha\left[\mathrm{PS}(Q) - \mathrm{PS}(Q_0)\right] > 0.$$

I proceed by bounding $\mathrm{val}(Q)$. By strong duality, $\mathrm{val}(Q)$ is equal to the value of the policymaker's dual problem, which can be written as (cf. Chapter 1.2 of Lasserre, 2009)

$$\min_{\lambda, \mu \in \mathbb{R}} \left\{\mu + \lambda\left[Q - \mathbf{E}[q_0(\theta, Q)]\right]\right\}$$
$$\text{s.t. } \mu \geq \frac{\int_{G^{-1}(1-s)}^{\overline{\eta}} \left[\eta - c(\delta) - \kappa - \delta p(Q) + \frac{\int_{\eta}^{\overline{\eta}} [\mathbf{E}[\omega \,|\, r] - 1] \,\, \mathrm{d}G(r)}{g(\eta)} - \lambda\left[\delta - \mathbf{E}[q_0(\theta, Q) \,|\, \eta]\right]\right] \, \mathrm{d}G(\eta)}{s}$$
$$\forall \; s \in (0, 1].$$

Suppose that $\delta < Q < Q_0$, so that $\delta - \mathbf{E}[q_0(\theta, Q)] < \delta - Q < 0$. By taking $s = 1$, notice that the constraint implies that

$$\lambda \leq -\frac{\mathbf{E}[\omega\eta] - c(\delta) - \kappa - \delta p(Q) - \mu}{\mathbf{E}[q_0(\theta, Q)] - \delta}.$$

Here, Fubini's theorem implies that

$$\int_{\underline{\eta}}^{\overline{\eta}} \left[\eta + \frac{\int_{\eta}^{\overline{\eta}} \left[\mathbf{E}[\omega \,|\, s] - 1\right] \,\, \mathrm{d}G(s)}{g(\eta)}\right] \, \mathrm{d}G(\eta) = \int_{\underline{\eta}}^{\overline{\eta}} \left[\eta + \int_{\underline{\eta}}^{\eta} \left[\mathbf{E}[\omega \,|\, \eta] - 1\right] \,\, \mathrm{d}s\right] \, \mathrm{d}G(\eta)$$
$$= \int_{\underline{\eta}}^{\overline{\eta}} \left[\eta + \mathbf{E}\left[\omega\left(\eta - \underline{\eta}\right) \,|\, \eta\right] - \left(\eta - \underline{\eta}\right)\right] \, \mathrm{d}G(\eta) = \mathbf{E}[\omega\eta].$$

Since $Q - \mathbf{E}[q_0(\theta, Q)] < 0$, it follows that

$$\begin{aligned}
\text{val}(Q) &\geq \mu + \frac{[\mathbf{E}[\omega\eta] - c(\delta) - \kappa - \delta p(Q) - \mu]\,[\mathbf{E}[q_0(\theta, Q)] - Q]}{\mathbf{E}[q_0(\theta, Q)] - \delta} \\
&= \frac{[\mathbf{E}[\omega\eta] - c(\delta) - \kappa - \delta p(Q)]\,[\mathbf{E}[q_0(\theta, Q)] - Q]}{\mathbf{E}[q_0(\theta, Q)] - \delta} + \frac{Q - \delta}{\mathbf{E}[q_0(\theta, Q)] - \delta} \cdot \mu.
\end{aligned}$$

Now, by taking $s \to 0$ and using L'Hôpital's rule, notice that the constraint in the policymaker's dual problem implies that

$$\mu \geq -\kappa.$$

It follows that

$$\begin{aligned}
\text{val}(Q) &\geq \frac{[\mathbf{E}[\omega\eta] - c(\delta) - \delta p(Q)]\,[\mathbf{E}[q_0(\theta, Q)] - Q]}{\mathbf{E}[q_0(\theta, Q)] - \delta} - \kappa \\
&= \frac{[\mathbf{E}\,[\omega\,[u(\delta, \theta) - v_0(\theta, Q)]] - c(\delta) - \delta p(Q)]\,[\mathbf{E}[q_0(\theta, Q)] - Q]}{\mathbf{E}[q_0(\theta, Q)] - \delta} - \kappa.
\end{aligned}$$

Consequently, a public option is optimal if

$$\begin{aligned}
&\mathbf{E}\,[\omega\,[v_0(\theta, Q) - v_0(\theta, Q_0)]] + \mathbf{E}\,[\omega\,[u(\delta, \theta) - v_0(\theta, Q)]] \cdot \frac{\mathbf{E}[q_0(\theta, Q)] - Q}{\mathbf{E}[q_0(\theta, Q)] - \delta} \\
&\quad > \kappa + [c(\delta) + \delta p(Q)] \cdot \frac{\mathbf{E}[q_0(\theta, Q)] - Q}{\mathbf{E}[q_0(\theta, Q)] - \delta} + \alpha\,[\text{PS}(Q_0) - \text{PS}(Q)]\,.
\end{aligned}$$

This concludes the proof of Proposition 13. □

Proposition 13 establishes a different sufficient condition for a public option to be optimal that can depend on global—rather than only local—properties of input supply. This is motivated by the fact that input supply can be flat in a neighborhood of the laissez-faire aggregate quality level Q_0, yet input supply might be inelastic outside of that neighborhood. In this case, the sufficient condition in Proposition 12 might not be satisfied as it only exploits information about input supply at Q_0. By contrast, the sufficient condition in Proposition 13 can exploit information about input supply away from Q_0.