Intermediated Implementation

Anqi Li* Yiqing Xing†

First Draft: December 2015
This Draft: October 2016

Abstract

Many real-world problems like sales, taxation and health care regulation, feature a principal, one or more intermediaries and agents with hidden characteristics. In these problems, intermediaries can specify the full menu of the multi-faceted consumption bundles that they offer to agents, whereas the principal is limited to regulating some but not all aspects of the bundles that agents consume, due to legal, information and administrative barriers. We examine how the principal can implement in these situations any target social choice rule that is incentive compatible, individually rational and feasible among agents. We show that when intermediaries have private values and are perfectly competitive, the principal's goal can be achieved by imposing a per-unit fee schedule that allows intermediaries to break even under the target social choice rule. When intermediaries have interdependent values or market power, per-unit fee schedule cannot generally be used to achieve the principal's goal, while regulating the distribution of limited aspects of sold bundles can. We study the policy implication of these results.

*Department of Economics, Washington University in St. Louis, anqili@wustl.edu.

†Carey Business School, Johns Hopkins University, xingyq@gmail.com. We thank Aaron Bodoh-Creed, Dirk Bergemann, Chen Cheng, Matt Jackson, Jon Levin, Xing Li, Justin Yifu Lin, George Mailath, Mike Peters, Andy Postlewaite, Florian Scheuer, Ilya Segal, Chris Shannon, Hugo Sonnenschein, Takuro Yamashita, as well as seminar participants at U of Michigan, UPenn, Stanford, ASSA, ES Summer Meetings and SED for helpful comments and suggestions. Xiaoyu Liu provided excellent graphical assistance. All errors are our own.

1 Introduction

Many real-world problems, such as sales, taxation and health care regulation, feature the interactions between a principal, one or more intermediaries and agents with hidden characteristics. In these problems, intermediaries can specify the full menu of the multi-faceted consumption bundles that they offer to agents, whereas the principal is limited to regulating some but not all aspects of the bundles that agents consume, due to legal, information and administrative barriers. Examples that fit this description abound. For instance, while a manufacturer controls the qualities of sold products, she cannot influence the prices that retailers charge to customers with hidden tastes.[1] While the government observes the wages earned by workers with hidden innate abilities, she cannot enforce other terms of the employment contract that is signed between firms and workers, such as the required performance for each level of pay. In countries that adopt market-based health care systems, many government policies, such as coverage subsidies, regulate some but not all aspects the insurance products that private companies sell to patients with hidden risk types.

This paper takes a step towards understanding what outcomes can be achieved in these settings when the principal and intermediaries have different objectives. Specifically, we fix any target social choice rule that the principal aims to implement in case she herself can dictate the full menu of bundles that is offered to agents, subject to agents' incentive constraints and feasibility constraints. We consider a game where intermediaries propose menus of consumption bundles to agents, taking as given the principal's partial regulation over limited aspects of sold bundles. We examine policies that achieve *intermediated implementation*, meaning that there exists a sub-game perfect equilibrium of the above described game where each agent consumes his bundle under the target social choice rule.

Policies that achieve intermediated implementation differ, depending on whether intermediaries have *private values* or *interdependent values*, and whether the intermediary industry is *competitive* or *monopolistic*. Specifically, intermediaries have private values if their payoffs do not depend directly on the agent's hidden characteristic, and they are competitive if more than one of them play the above described game. We show that if intermediaries have private values and are competitive, then

[1] In 1980, the Sherman Act's complete ban of vertical price fixing became effective again. Since then, retail price maintenance has been illegal in the U.S..

intermediated implementation can be achieved by a per-unit fee schedule that leaves intermediaries zero profit from selling the target bundle of any agent type. However, if intermediaries have interdependent values or monopoly power, then per-unit fee schedules cannot generally be used to achieve intermediated implementation. In these situations, regulating the distribution of limited aspects of sold bundles fulfills the principal's purpose.

To illustrate these results, consider the problem between a car manufacturer, one or more dealers, and a continuum of customers whose hidden tastes for car quality are either high or low. In case the manufacturer controls both the quality and the price of cars, she can simply offer a menu of quality-price bundles and let agents self-select. But in reality, it is dealers who have the freedom to set car prices, whereas the manufacturer observes the quality of each sold car and charges dealers a fee accordingly. Consider two commonly used fee schedules: an invoice-price schedule that charges different per-unit fees for the various kinds of sold cars, as well as a quota scheme that requires dealers to sell a certain variety of cars.[2] We examine which fee schedules implement the target quality-price rule, which maximizes the manufacturer's objective in case she controls both the quality and the price, subject to customers' incentive constraints and feasibility constraints.

Our first result shows that if the profit from car sales does not depend directly on the customer's taste and if dealers are competitive, then the manufacturer's goal can be achieved by the invoice-price schedule that leaves dealers zero profit from selling the target bundle of any customer type. This result exploits only the assumptions that intermediaries have private values and are competitive, as well as the fact that the target quality-price rule is incentive compatible for customers. Specifically, under the current invoice-price schedule, a bundle makes a profit if and only if it charges a higher-than-target-level price. Moreover, this is true regardless of the customer's type, which has no direct effect on sales profit under the assumption of private values. Meanwhile, since the target quality-price rule is incentive compatible for customers, it follows that each customer most prefers his target bundle among all target bundles, let along those bundles that are profitable to dealers. Thus, the competition between dealers drives profits to zero and sustains the target quality-price rule in every subgame perfect equilibrium of the contracting game.

[2]A recent quote from a salesman at a Benz dealership in the Bay Area goes: "*we have four 2015 CLA-class left and they must go, otherwise Benz won't deliver us any new cars.*"

When intermediaries have interdependent values or market power, we can no longer use per-unit fee schedules to implement the target social choice rule. In the case of interdependent values, this negative result generalizes the insight of Rothschild and Stiglitz (1976): in our example, if the profit from car sales depends directly on the customer's type (e.g., high-taste customers are difficult to serve), then dealers can make a profit by offering pooling bundles that incur low invoice-prices to multiple types of customers. In the case of market power, this negative result follows from the logic of monopolistic screening: under most invoice-price schedules of interest, dealers can distort the prices paid by some customers in order to extract information rents from other customers.

In these situations, regulating the distribution of limited aspects of sold bundles helps achieve intermediated implementation. In our example, suppose that 20 percent of customers have a high taste for quality and the remaining 80 percent have a low taste for quality. Consider a policy that requires every dealer to sell 20 percent of high quality cars and 80 percent of low quality cars. In order to meet this requirement, the only way that dealers can deviate from serving the target quality car to each customer is to sell high quality cars to low taste customers, and low quality cars to high taste customers. However, since this permuted quality rule is decreasing in the customer's taste, it is not *implementable*, meaning that it cannot be part of any incentive compatible social choice rule and thus cannot arise in any equilibrium outcome. Thus in equilibrium, dealers must offer the target quality car to each customer. Based on this result, we then use either the competition between dealers or the envelope theorem to enforce the target levels of prices.

In the general case where a bundle consists of a consumption good and a price, we characterize target consumption rules that are not implementable after any permutation. By exploiting intrinsic properties of implementable consumption rules, such as monotonicity or cyclic monotonicity (Rochet (1987)), we show that permuting target consumption goods across agents does not yield implementable consumption rules, if such permutation changes the total utility of consumption among agents (henceforth (DU)). Furthermore, we show that (DU) can be easily satisfied if (1) the type space and consumption space are single-dimensional, and agent utilities satisfy the single-crossing property, or if (2) the type space is multi-dimensional, the consumption space is rich, and agent utilities vary significantly with the consumption.

We give several applications of these results. First, we consider how the government can use wage regulations to achieve redistribution among workers with hidden innate abilities, leaving firms to specify and to enforce other terms of the employment contract such as the required performance for each level of pay. In this application, we show that Mirrlees (1971)'s income tax schedule still attains constrained efficiency, if firms do not benefit directly from workers' innate abilities and if the labor market is competitive. When any of these conditions fails, income taxation per se cannot be used to attain many redistributive goals, whereas a policy that calls firms to match the target wage distribution restores the effectiveness of the Mirrleesian tax schedule. In reality, this second policy resembles employment targets that mandate employers fill a certain percentage of their positions with specific kinds of people, such as people with hidden disabilities.

We finally consider how the government can induce private companies to provide the desired amount of health insurance through coverage plan subsidies. This application features interdependent values, as the profit from insurance sales depends directly on the patient's risk type. Thus in general, per-unit coverage subsidies cannot induce the desired amount of insurance provision, whereas regulating the variety of sold coverage plans can. This result justifies the policies that the Affordable Care Act recently introduced, which mandate that all participating companies in the new health insurance exchange offer a variety of coverage plans for different patient types,[3] and penalize companies for selling too many low coverage plans.[4]

1.1 Related Literature

Private value vs. interdependent value Whether intermediaries have private values and interdependent values hinges on the contracting technology that parties use. To illustrate, recall that in the literature on non-linear taxation, a commonly made assumption is that firms benefit from the worker's output, which in turn depends on the worker's innate ability and labor hours. Thus if output is contractible as in Mirrlees (1971), then firms have private values because their profit does not depend directly on the worker's ability. However, if labor hours are contractible as in Stantcheva (2014), then firms have interdependent values because their profit depends

[3] Jean Folger, "How to Choose Between Bronze, Silver, Gold and Platinum Health Insurance Plans," *Forbes*, October 1, 2013.

[4] "Explaining Health Care Reform: Risk Adjustment, Reinsurance, and Risk Corridors," *The Kaiser Family Foundation*, January 22, 2014.

directly on the worker's ability.

Competitive equilibria with adverse selection As Rothschild and Stiglitz (1976) shows, interdependent values, or adverse selection, may result in the non-existence of competitive equilibria in a decentralized insurance market. In order to overcome this challenge, subsequent studies have imposed various restrictions on how economic agents behave in competitive environments. For example, Prescott and Townsend (1984), Bisin and Gottardi (2006) and many others examine the outcome of price-taking behaviors in a general equilibrium economy. Miyazaki (1977), Wilson (1977), Riley (1979), Netzer and Scheuer (2014) and the references therein adopt equilibrium concepts that allow players to anticipate the reaction of other players to their contract offers.[5] And Guerrieri et al. (2010) demonstrates how search and match frictions help restore equilibrium existence. Among these studies, Rothschild and Stiglitz (1976), Wilson (1977) and Riley (1979) limit players to offering one single contract.

The current paper differs from these studies in several aspects. First and foremost, we examine policy interventions that implement through intermediaries any social choice rule that is incentive compatible, individually rational and feasible among agents, rather than to characterizing the outcome of a laissez-faire economy. Second, we allow intermediaries to specify all aspects of the consumption bundle, impose no limit on the size of the menu, and consider a standard game where intermediaries offer menus of consumption bundles to agents and let agents self-select, which all make it harder, not easier, to establish equilibrium existence. Stantcheva (2014) studies the optimal redistribution policy in the presence of decentralized employment contracting. However, she adopts Miyazaki (1977)'s equilibrium concept, and this drives the key difference between our results.

Common agency game Since the contracting game between intermediaries and agents is a common agency game, it follows from the existing literature on common agency games (see, e.g., Epstein and Peters (1999), Martimort and Stole (2002)) that in the current setting, the revelation principle breaks down, and restricting intermediaries to offering menus of consumption bundles is not without loss. Nevertheless, even if general contracts are allowed, our results remain valid because under our policies, all intermediaries offering the full menu of target bundles remains an equilibrium of the contracting game.

[5] Some but not all of these studies provide strategic foundations for these equilibrium concepts.

Regulating distribution Jackson and Sonnenschein (2007) (henceforth JS) shows that in a screening model where agent types are i.i.d. across many replicas of the same decision problem, we can virtually implement Pareto-efficient allocations by requiring the distribution of reported types to match the distribution of true types. The current analysis differs from JS in three aspects. First, the setting of JS involves no intermediary whereas we examine the possibility of achieving intermediated implementation. Second, JS regulates the distribution of reported types whereas we regulate the distribution of limited aspects of sold bundles. Third, JS's result hinges on the Pareto efficiency of the target allocation, whereas our construction exploits the cyclic monotonicity of implementable consumption rules.

Rochet (1987) shows that in screening models where the agent has quasi-linear utility and multi-dimensional hidden characteristics, a consumption rule is implementable if and only if it satisfies cyclic monotonicity. Rahman (2011) points out that a consumption rule is cyclically monotone if and only if every permutation of it is weakly unprofitable for the agent. In our setting, distribution regulation achieves intermediated implementation if the target consumption rule is implementable whereas any permutation of it is not. This is equivalent to (DU), which requires that any permutation of the target consumption rule changes the total utility among agents.

Contracting with externality and moral hazard In our setting, each intermediary's contract offer has direct effect on the payoff of other intermediaries. In the meantime, our principal suffers from a moral hazard problem, as she is limited to regulating some but not all aspects of the bundles that are being sold. The first feature relates the current analysis to the literature on contracting with externalities (see, e.g., Segal (1999) and the references therein). The second feature is absent from Segal (1999), which examines a contracting problem with externalities where the principal dictates the full allocation among agents.

The remainder of this paper proceeds as follows: Section 2 introduces the model setup and presents leading examples; Section 3 states and analyzes the main results; Section 4 investigates several extensions; Section 5 concludes. See Appendix A for mathematical proofs, as well as the online Appendix for additional results.

2 Baseline Model

In this section, we first introduce the model setup and then demonstrate the applicability of our framework to sales, taxation and health-care regulation.

2.1 Setup

The economy consists of a principal, an integer number I of identical intermediaries and a unit mass of infinitesimal agents. Each agent independently draws a hidden characteristic θ from a measurable space (Θ, Σ) according to a probability function P_θ. A consumption bundle (x, y) consists of two elements x and y. The spaces of these elements, denoted by X and Y, are Borel subsets of $\mathbb{R}^d$ and $\mathbb{R}$, and are equipped with the restricted Borel σ-algebra $\mathcal{X}$ and $\mathcal{Y}$. Let Y be connected in $\mathbb{R}$. Each agent θ earns a utility $u(x, y, \theta)$ from consuming at most one bundle (x, y), whereas a constant-returns-to-scale technology yields a gross profit $\pi(x, y, \theta)$ from serving a bundle (x, y) to type θ agent.[6] Throughout, let u and π satisfy the following monotonicity condition:

Assumption 1. *$u(x, y, \theta)$ is decreasing in y and $\pi(x, y, \theta)$ is increasing in y.*

Let $\mathbf{0}$ denote the null bundle that yields a zero reservation payoff to all parties, i.e., $u(\mathbf{0}, \theta) = 0$ and $\pi(\mathbf{0}, \theta) = 0$ for all θ. Think of $\mathbf{0}$ as the outside option that parties pursue when they abstain from the market, and suppose without loss of generality that $\mathbf{0} \in X \times Y$.

Target social choice rule A social choice rule $(x, y) : \Theta \to X \times Y$ is a deterministic mapping between the agent's type space and the consumption bundle space. For now, suppose that the principal owns the technology for serving agents, and that her objective function f can depend on each agent's type and consumption bundle. The principal wants to implement a *target social choice rule* $(\hat{x}, \hat{y}) : \Theta \to X \times Y$ that maximizes her objective function, subject to agents' incentive compatibility and

[6] Appendix B.3 investigates an extension where the technology for serving agents can differ across intermediaries and exhibit non-constant returns to scale.

individual rationality constraints, as well as an unspecified feasibility constraint:

$$\begin{aligned}
&\max_{(x,y)(\cdot)} f\left\{(x(\theta), y(\theta), \theta) : \theta \in \Theta\right\}, \\
&\text{s.t. } u(x(\theta), y(\theta), \theta) \geq u\left(x(\theta'), y(\theta'), \theta\right), \forall \theta, \theta', && \text{(IC)} \\
&\text{and } u(x(\theta), y(\theta), \theta) \geq 0, && \text{(IR)} \\
&\text{and } F\left\{(x(\theta), y(\theta), \theta) : \theta \in \Theta\right\} \geq 0. && \text{(Feasibility)}
\end{aligned}$$

Denote by $\hat{x}(\Theta)$, $\hat{y}(\Theta)$ and $(\hat{x}, \hat{y})(\Theta)$ the image of Θ under the mapping $\hat{x}$, $\hat{y}$ and $(\hat{x}, \hat{y})$, respectively. Notice that for each $x \in \hat{x}(\Theta)$, there exists a unique $y \in \hat{y}(\Theta)$ (hereafter denoted by $\hat{y}(x)$) such that $(x, y) \in (\hat{x}, \hat{y})(\Theta)$. The reason is straightforward: if there exist $y, y' \in \hat{y}(\Theta)$ with $y < y'$ such that $(x, y), (x, y') \in (\hat{x}, \hat{y})(\Theta)$, then all agents will prefer (x, y) to (x, y') and hence $(x, y') \notin (\hat{x}, \hat{y})(\Theta)$, a contradiction. Define

$$\hat{\pi}(x) = \mathbf{E}_\theta\left[\pi\left(x, \hat{y}(x), \theta\right) \mid (\hat{x}, \hat{y})(\theta) = (x, \hat{y}(x))\right] \tag{2.1}$$

as the expected profit from serving a target consumption bundle $(x, \hat{y}(x))$ to its corresponding agents.

Intermediated implementation In the case where the principal can dictate the menu of consumption bundles (hereafter, *direct implementation*), she will simply offer the menu $\{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ of target consumption bundles and let agents self-select. Now suppose instead that the technology of serving agents is owned by intermediaries, whereas the principal — whose objective function f differs generally from π — can only regulate the x-dimension of sold bundles, due to the various legal, information and administrative barriers that we will further discuss in Section 2.2. The current paper concerns how the principal can implement the target social choice rule in this setting.

To formalize this problem, let μ_i be the measure on $(X \times Y \times \Theta, \mathcal{X} \otimes \mathcal{Y} \otimes \Sigma)$ that is induced by the bundles that intermediary i sells, and $\mu_{i,x}$ be the measure on $(X, \mathcal{X})$ that is induced by μ_i. Suppose that for each $i = 1, \cdots, I$, the principal observes only $\mu_{i,x}$, based on which she can charge a fee $\psi(\mu_{i,x})$ and leave intermediary i with the following net profit:

$$\int_{(x,y,\theta)} \pi(x, y, \theta) d\mu_i - \psi(\mu_{i,x}). \tag{2.2}$$

A fee schedule ψ achieves *intermediated implementation* if there exists a sub-game perfect equilibrium of the following game where each type θ agent consumes his target consumption bundle $(\hat{x}(\theta), \hat{y}(\theta))$:

1. The principal commits to ψ;
2. Intermediaries simultaneously propose a menu of deterministic consumption bundles $\sigma_i \subset 2^{X \times Y}$, $i = 1, \cdots, I$;
3. Each agent selects a bundle from $\sigma = \cup_{i=1}^{I} \sigma_i \cup \{\mathbf{0}\}$;
4. Intermediaries deliver selected bundles to agents and pay fees to the principal.

The above described contracting game best represents environments where intermediaries possess a rich strategy space (i.e., specify all aspects of the consumption bundle and faces no limitation on the size of the menu), whereas agents have no problem finding the bundle that they most prefer. In any sub-game perfect equilibrium of this game, each intermediary maximizes its net profit, taking as given the fee schedule, the strategies of other intermediaries and the fact that agents are utility maximizing.[7]

Two things are noteworthy. First, the payoff of any player in any equilibrium is bounded below by zero. Second, in any equilibrium that achieves intermediated implementation, every sold bundle (x, y) is known to satisfy $x \in \hat{x}(\Theta)$ and $y = \hat{y}(x)$ and to yield the target amount of profit $\hat{\pi}(x)$. By taxing this amount of profit away from the intermediary, the principal achieves the exact same outcome as in the case of direct implementation, where she confiscates the profit from serving the target consumption bundle to each type of agent.[8]

Policies Throughout this paper we mainly consider two kinds of policies. The first kind of policies, hereafter referred to as *per-unit fee schedules*, charges a fee $t(x)$

[7] We bestow intermediaries with a rich strategy space and adopt a standard solution concept in order to (1) best capture the rich strategic interactions between players in reality, and to (2) make it harder, not easier, to establish equilibrium existence in the case of interdependent values and perfect competition (see Section 3.1.2 for further details). See Section 1.1 for discussions of how related studies obtain equilibrium existence by imposing further restrictions on what intermediaries can offer (and providing strategic foundations for these restrictions) or by adding frictions to the contracting game.

[8] In case the principal's objective function f depends also on the intermediaries' profits, the target social choice rule prescribes a consumption bundle $(\hat{x}(\theta), \hat{y}(\theta))$ to each agent θ and a net profit $\hat{\Pi}_i$ to each intermediary i. In any equilibrium that achieves intermediated implementation, the principal first taxes away the profit from serving agents and then makes a lump sum transfer $\hat{\Pi}_i$ to each intermediary i in order to achieve the exact same outcome as in the case of direct implementation.

for every sold bundle (x, y) whereby $x \in \hat{x}(\Theta)$, as well as a big penalty (denoted by "$+\infty$", meaning that the penalty deters this bundle from being offered in equilibrium) for every sold bundle (x, y) whereby $x \notin \hat{x}(\Theta)$.[9] Under a per-unit fee schedule, the total charge to intermediary i is equal to

$$\psi_{per-unit}\left(\mu_{i,x}\right) = \int_{x \in \hat{x}(\Theta)} t(x) d\mu_{i,x}. \tag{2.3}$$

The second kind of policies imposes aggregate restrictions on the x-dimension of the bundles that are finally sold. Much attention will be given to *distribution regulations* ψ_{distr} which implements a fee schedule ψ if the intermediary matches the target probability distribution over x, and inflict a big penalty on the intermediary otherwise. Formally, let $\hat{P}_x$ denote the probability measure on $(X, \mathcal{X})$ that is induced by the target social choice rule, where $\hat{P}_x(x \in A) = P_\theta\left(\theta : \hat{x}(\theta) \in A\right)$ for each $A \in \mathcal{X}$. Then the policy of our interest can be expressed as follows:

$$\psi_{distr}\left(\mu_{i,x}\right) = \begin{cases} \psi\left(\mu_{i,x}\right) & \text{if } \frac{\mu_{i,x}}{\int_{x \in \hat{x}(\Theta)} d\mu_{i,x}} = \hat{P}_x, \\ +\infty & \text{otherwise.} \end{cases} \tag{2.4}$$

Key Conditions The current analysis hinges on two key conditions: whether intermediaries have *private values* or *interdependent values*, and whether the intermediary industry is *perfectly competitive* or *monopolistic*. Formally, intermediaries have private values if their payoff does not depend directly on the agent's hidden type, and they have interdependent values otherwise:

Definition 1. *Intermediaries have private values if $\pi(x, y, \theta)$ is independent of θ for all $(x, y) \in X \times Y$, and they have interdependent values if $\pi(x, y, \theta)$ depends on θ for some $(x, y) \in X \times Y$.*

Meanwhile, the intermediary industry is said to be perfectly competitive if the above described contracting game involves multiple intermediaries, and it is said to be monopolistic otherwise:

[9] Big penalties are commonly used to eliminate unwanted outcomes in mechanism design. As an example, recall the taxation principle which says that in the case of direct implementation, every menu $\{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ of consumption-price bundles can be rewritten as a tax schedule $T(x) = \begin{cases} \hat{y}(x) & \text{if } x \in \hat{x}(\Theta); \\ +\infty & \text{if } x \notin \hat{x}(\Theta). \end{cases}$

Definition 2. *The intermediary industry is perfectly competitive if $I \geq 2$, and it is monopolistic if $I = 1$.*

2.2 Examples

2.2.1 Intermediated Sales

A manufacturer, say Mercedes Benz, faces many customers whose private taste is denoted by θ. Each customer has a single-unit demand for the various car models $X = \{\text{CLA-class}, \text{C-class}, \text{E-class}, \text{ etc.}\}$ that the manufacturer offers. A consumption bundle (x, y) specifies a car model $x \in X$ and a price $y \in Y = \mathbb{R}_+$; it incurs a manufacturing cost $c^m(x)$ and yields a payoff $u(x, y, \theta) = v(x, \theta) - y$ to type θ customer. The gross profit from distributing (x, y) to type θ customer is $\pi(x, y, \theta) = y - c^d(x, \theta)$, where $c^d(x, \theta)$ denotes the distributional cost. In principle, $c^d(x, \theta)$ can depend on θ because high-taste customers can be more difficult to serve than low-taste customers.

In the case where the manufacturer owns both the manufacturing technology and the distribution technology, her net profit from serving a bundle (x, y) to type θ agent is given by $\pi(x, y, \theta) - c^m(x)$. The target social choice rule maximizes the manufacturer's expected profit, subject to customers' incentive compatibility and individual rationality constraints, i.e., $\max_{(x,y)(\cdot)} \int \pi(x(\theta), y(\theta), \theta) - c^m(x(\theta)) dP_\theta$ s.t. (IC) and (IR).

In reality, the distribution technology is owned by dealers, who are key for the manufacturer to get access to its customers. For simplicity, suppose there are I identical dealers who independently decide the price of each car model. By law (see Footnote 1 for Sherman's Act's ban on vertical price fixing), the manufacturer cannot exert influence on dealers' pricing strategies or force dealers to disclose their charged prices or earned profits. Nevertheless, the manufacturer observes the quantity of each sold car model, based on which she charges a fee $\psi(\mu_{i,x})$ and leaves dealer i with a net profit $\int_{(x,y,\theta)} \pi(x, y, \theta) d\mu_i - \psi(\mu_{i,x})$.

In the current example, whether dealers have private or interdependent values depends on whether the distribution cost $c^d(x, \theta)$ varies with the customer's hidden taste or not. The contracting game evolves as follows: first, the manufacturer commits to a fee schedule; second, dealers simultaneously propose menus of car model-price bundles; finally, customers purchase their most preferred bundle and dealers pay fees to the manufacturer. The policies of our interest are commonly used in reality:

- $\psi_{per-unit}$ represents *invoice-price schedules* that charge a distinct per-unit fee for each car model.

- ψ_{distr} regulates the variety of sold cars. For example, Mercedes Benz can refuse to supply more C-class cars until a certain number of CLA-class cars have been sold. A recent quote from a salesman at a Benz dealership in the Bay Area goes: "*we have four 2015 CLA-class left and they must go, otherwise Benz won't deliver us any new cars.*"

2.2.2 Redistribution with Decentralized Employment Contracting

A government faces a continuum of workers whose hidden innate ability is denoted by θ. An employment contract (x, y) specifies the worker's after-tax consumption $x \in X \subset \mathbb{R}$ and required performance $y \in Y \subset \mathbb{R}_+$; it yields a utility $u(x, y, \theta) = x - v(y, \theta)$ to type θ worker and adds a value $\pi(x, y, \theta) = h(y, \theta) - x$ to his employer and the society.

The government wants to implement the social choice rule $(\hat{x}, \hat{y})$ that maximizes the weighted sum of worker utilities $\int \lambda(\theta)(x(\theta) - v(y(\theta), \theta))dP_\theta$ subject to (IC), (IR) and a resource constraint $\int h(y(\theta), \theta) - x(\theta)dP_\theta \geq R$. In a centralized economy where the government dictates the menu of employment contracts, she will simply propose the menu $\{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ and let workers self-select. After firms passively execute the signed contracts, each worker surrenders an income tax $\hat{\pi}(x)$ to the government, retains a consumption x to himself and leaves the firm with a zero profit. This is the famous income taxation problem studied by Mirrlees (1971).[10]

But in market economies, firms have the freedom to specify what kind of performance is required for each level of pay. Furthermore, in large and complex organizations, worker performance is best observed by his employer, and cannot be easily inferred by the government based on records that depend on (random) confounding factors (e.g., the company's stock price). Formally, let I identical firms independently decide the required performance for each level of after-tax consumption. The government observes the consumption of each worker hired by each firm i, based on which she charges a tax $\psi(\mu_{i,x})$ and leaves firm i a with net profit $\int_{(x,y,\theta)} \pi(x, y, \theta)d\mu_i - \psi(\mu_{i,x})$.

[10] Alternatively, let an employment contract $(x + \hat{\pi}(x), y)$ consist of a pre-tax income $x + \hat{\pi}(x)$ and a required performance level y. In order to map each level x of after-tax consumption to a unique level $x + \hat{\pi}(x)$ of pre-tax labor income, we shall henceforth require that $x/\hat{\pi}(x)$ differ across x.

In the current example, whether firms have private values or interdependent values depends on the contracting technology: if y represents the worker's effective labor output as in Mirrlees (1971), then firms have private values because $h(y, \theta) = y$ is independent of θ; however, if y represents labor hours as in Stantcheva (2014), then firms have interdependent values because $h(y, \theta)$ depends directly on θ (e.g., $h(y, \theta) = \theta y$ in Stantcheva (2014)).[11] The contracting game evolves as follows: first, the government commits to a policy; second, firms simultaneously propose menus of employment contracts, and workers opt into the contract that they most prefer; finally, contracts are executed, and the government collects taxes from firms. The policies of our interest are commonly used in reality:

- $\psi_{per-unit}$ represents income tax schedules that charge a per-worker tax $t(x)$ for each level of x. The Mirrleesian tax schedule is a special case where $t(x) = \hat{\pi}(x)$ for each x, meaning that it is designed purposefully to make firms break even under the target social choice rule.

- ψ_{distr} requires that firms achieve the target wage distribution among their employees. In reality, this policy resembles employment targets which mandate that big employers fill a certain percentage of their positions with, for example, people with (hidden) disabilities.

2.2.3 Market-Based Health Care Regulation

A government faces a continuum of patients whose hidden risk type $\theta = (\theta_1, \cdots, \theta_d)$ is randomly drawn from $\Theta = \Delta^d$. Type θ patient's endowment is equal to $e_s \in \mathbb{R}$ with probability θ_s, $s = 1, \cdots, d$. An insurance product (x, y) consists of a state-contingent consumption plan $x = (x_1, \cdots, x_d) \in \mathbb{R}^d$ and a premium $y \in \mathbb{R}$; it yields an expected utility $u(x, y, \theta) = \sum_{s=1}^{d} \theta_s \cdot v(x_s - y)$ to type θ patient, as well as an expected profit $\pi(x, y, \theta) = y - \sum_{s=1}^{d} \theta_s \cdot (x_s - e_s)$.

In the case where the government administers a single-payer health care system, she would offer the menu of insurance products that maximizes the weighted sum of patient utilities $\int \lambda(\theta) u(x(\theta), y(\theta), \theta) dP_\theta$ subject to (IC) and (IR) and a budget constraint $\int \pi(x(\theta), y(\theta), \theta) dP_\theta \geq 0$. However, in countries like the U.S. that adopt market-based health care systems, insurance policies are sold by private companies, while many governmental policies regulate some but not all aspects of the sold policies.

[11] Section 4 discusses the policy implication of this subtle observation.

For concreteness, we mainly consider consumption-plan regulations (e.g., coverage subsidy) in the main body of this paper, though the same analysis can be applied to the study of premium regulations (e.g., premium subsidy).

Formally, let $X = \{\text{platinum}, \text{gold}, \text{silver}, \text{bronze}, \text{etc.}\} \subset \mathbb{R}^d$ be the space of coverage plans and $Y \subset \mathbb{R}$ be the space of premiums. For simplicity, suppose there are I homogeneous insurance companies who independently decide the price of each coverage plan. The government observes the type of each sold coverage plan, based on which she charges a fee $\psi(\mu_{i,x})$ and leaves insurance company i with a net profit $\int_{(x,y,\theta)} \pi(x, y, \theta) - \psi(\mu_{i,x})$.

In the current example, the fact that $\pi(x, y, \theta)$ depends directly on θ means that insurance companies have interdependent values. Time evolves as follows: first, the government commits to a fee schedule; second, insurance companies simultaneously propose insurance policies; finally, patients purchase their most preferred policies and money changes hands between insurance companies and the government. The policies of our interest have real-life counterparts:

- $\psi_{per-unit}$ represents per-unit coverage subsidies;
- ψ_{distr} requires insurance companies to sell a variety of coverage plans. Recently similar rulings have been introduced by the Affordable Care Act, which mandate that all participating companies in the new health insurance exchange offer a variety of plans for different patient types, and penalize companies for selling too many low coverage plans.

3 Main Results

The main research addressed by the current paper concerns which policies of interest can achieve intermediated implementation under what conditions. Table 1 summarizes the main results. The message is twofold. First, in the case where intermediaries have private values and are perfectly competitive, intermediated implementation can be achieved by imposing a per-unit fee schedule that charges $t(x) = \hat{\pi}(x)$ for every $x \in \hat{x}(\Theta)$. This fee schedule is interesting for two reasons: first, it is the exact tariff schedule that the *taxation principle* prescribes in the case of direct implementation; second, it leaves intermediaries with a zero net profit if every $x \in \hat{x}(\Theta)$ is "correctly priced" at $\hat{y}(x)$. Under this fee schedule, the total charge to intermediary i is equal

Table 1: Which policies of our interest achieve intermediated implementation.

Panel A: per-unit fee

	competitive	monopolistic
private value	yes	no
interdependent value	no	no

Panel B: distribution regulation

	competitive	monopolistic
private value	yes	yes
interdependent value	yes	yes

to

$$\hat{\psi}_{per-unit}\left(\mu_{i,x}\right) = \int_{x\in\hat{x}(\Theta)} \hat{\pi}(x) d\mu_{i,x}. \tag{3.1}$$

Second, if intermediaries have interdependent values or market power, then in general, we cannot use per-unit fee schedules — which include but are not limited to $\hat{\psi}_{per-unit}$ — to implement the target social choice rule. In these situations, the following distribution regulation — which charges $\hat{\psi}_{per-unit}$ if and only if the intermediary matches the probability distribution over x under the target social choice rule — fulfils our purpose:

$$\hat{\psi}_{distr}\left(\mu_{i,x}\right) = \begin{cases} \hat{\psi}_{per-unit}\left(\mu_{i,x}\right) & \text{if } \frac{\mu_{i,x}}{\int_{x\in\hat{x}(\Theta)} d\mu_{i,x}} = \hat{P}_x, \\ +\infty & \text{otherwise.} \end{cases} \tag{3.2}$$

The remainder of this paper is devoted to analyzing each cell of these tables.

3.1 Per-Unit Fee Schedule

This section examines the effectiveness of per-unit fee schedules in various environments.

3.1.1 Private Value and Perfect Competition

Our first theorem shows that in the case where intermediaries have private values and are perfectly competitive, intermediated implementation can be achieved by imposing the per-unit fee schedule $\hat{\psi}_{per-unit}$ that allows intermediaries to break even under the target social choice rule.

Theorem 1. *Suppose that intermediaries have private values and are perfectly competitive, and that Assumption 1 holds. Then under $\hat{\psi}_{per-unit}$, the contracting game has a sub-game perfect equilibrium where $\sigma_i^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}, i = 1, \cdots, I$.*

Theorem 1 says that when the customer's hidden taste has no direct effect on the distribution cost and dealers are competitive, a manufacturer can achieve her goal through a simple invoice-price schedule even if she has no direct control over the dealers' pricing strategies. When worker ability has no direct effect on a firm's profit and the labor market is competitive, the government can attain constrained efficiency through the Mirrleesian tax schedule even if firms can freely decide the required performance for each level of pay.

To illustrate this result, suppose that $\hat{x}(\Theta) = X$ and $\hat{y}(\Theta) = Y$, and that X and Y are both connected in $\mathbb{R}$. Figure 1 depicts the net profit of all feasible consumption bundles under $\hat{\psi}_{per-unit}$. Since $\hat{\psi}_{per-unit}$ allows intermediaries to break even from serving the target consumption bundle to any type of agent, it follows that all bundles in region Π^+ make a profit under $\hat{\psi}_{per-unit}$. In particular, this result is independent of the agent's type because of the assumption of private values. At first sight, it is not obvious what prevents the bundles in region Π^+ from being offered or accepted in any equilibrium outcome.

Figure 2 adds the agents' indifference curves to Figure 1. Since the target social choice rule is incentive compatible, it follows that the graph $\{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ of target consumption bundles constitutes a lower envelope of region U^+ where each agent obtains at least his utility under the target social choice rule. Since regions U^+ and Π^+ are fully separated, it follows that no intermediary can make a profit when the other intermediaries are offering the full menu of target consumption bundles. Thus all intermediaries offering the full menu of target consumption bundles constitutes an equilibrium outcome of the contracting game.

Alternative interpretation Theorem 1 is closely related to welfare theorems and classical results on price competition. Specifically, consider an augmented economy

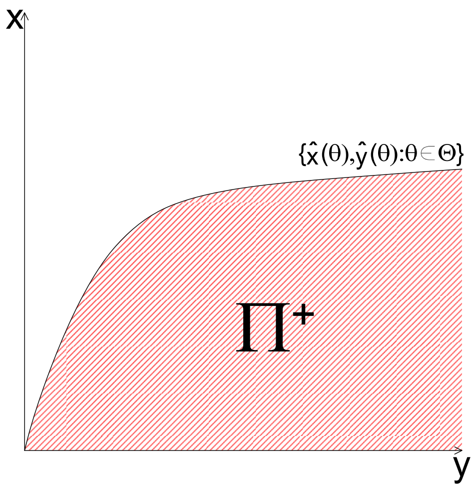

Figure 1: Net profit under $\hat{\psi}_{per-unit}$.

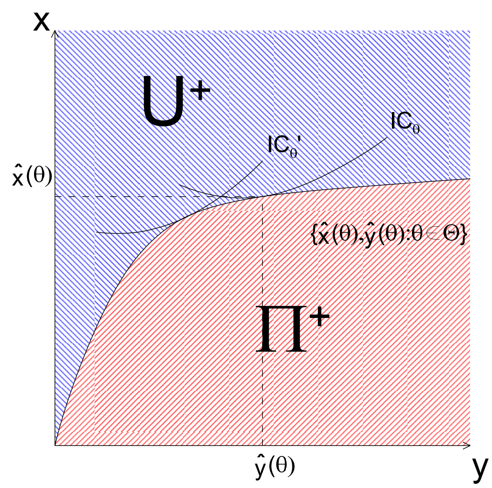

Figure 2: Equilibrium analysis.

where the graph of target consumption bundles represents the frontier of the intermediary's production technology, and each agent most prefers his target consumption bundle among all bundles lying on this frontier. In this economy, the target social choice rule is Pareto-efficient and can thus be sustained by the Bertrand competition between intermediaries.

Irrelevance of contracting protocol Theorem 1 is independent of the contracting protocol that parties use. In the opposite situation where agents make take-it-or-leave-it offers to intermediaries, the contracting game induced by $\hat{\psi}_{per-unit}$ has an equilibrium where every agent proposes his target consumption bundle in order to maximize his payoff while allowing intermediaries to break even. As the reader will soon realize, this neutrality result breaks down when intermediaries have interdependent values.

Outcome uniqueness The next corollary shows that under mild regularity conditions, all agents obtain their utility under the target social choice rule and all intermediaries break even in every sub-game perfect equilibrium of the contracting game.

Corollary 1. *Suppose that intermediaries have private values and are perfectly competitive, that Assumption 1 holds, and that* $\pi(x, y)$ *is continuous in* y *and* $\inf_{y \in Y} \pi(x, y) \leq 0$ *for every* x*. Then under* $\hat{\psi}_{per-unit}$*, all agents obtain their utility under the target* *social choice rule and all intermediaries break even in every sub-game perfect equilibrium of the contracting game.*

To better understand this result, notice first that in any equilibrium of the contracting game, every sold bundle must yield zero profit because otherwise some intermediary can undercut and make a profit. Second, if any type θ agent's equilibrium utility is below its counterpart under the target social choice rule, then any intermediary can add $(\hat{x}(\theta), \hat{y}(\theta) + \epsilon)$ to his menu and make a profit. Combining these observations yields the result.

General contract space Since our contracting game is a common agency game with privately informed agents, it follows from the existing literature on common agency games (see, e.g., Epstein and Peters (1999), Martimort and Stole (2002)) that in the current setting, the revelation principle breaks down and restricting intermediaries to offering menus of consumption bundles is not without loss. However, even if general contracts are allowed, Theorem 1 remains valid because under $\hat{\psi}_{per-unit}$, all intermediaries offering the full menu of target consumption bundles remains an equilibrium of the contracting game.

Value-added tax In case the principal can regulate the intermediary's profit, too, intermediated implementation can be easily achieved by imposing a hundred percent value-added tax. But in reality, this solution may not work as well as it seems for two reasons. First, non-governmental entities (e.g., manufacturer) cannot force intermediaries to disclose earned profits. Second, a hundred percent value-added tax is seldom observed in reality. In the case where intermediaries already pay a 100α percent value-added tax, we can still achieve intermediated implementation by charging a modified per-unit fee $t(x) = (1-\alpha)\hat{\pi}(x)$ for every sold bundle that satisfies $x \in \hat{x}(\Theta)$.

3.1.2 Interdependent Value and Perfect Competition

In this section, we explain why per-unit fee schedules may fail to achieve intermediated implementation in the case where intermediaries have interdependent values and are perfectly competitive. We begin with a simple example:

Example 1. *Suppose $X \subset \mathbb{R}$, $\Theta = \left\{\overline{\theta}, \underline{\theta}\right\} \subset \mathbb{R}$ and $I \geq 2$. Agent utility $u(x, y, \theta) = v(x, \theta) - y$ is quasi-linear in y and satisfies the single-crossing property (hereafter*

SCP) in (x, θ). $\pi(x, y, \theta)$ is continuous in y and satisfies $\pi\left(x, y, \overline{\theta}\right) > \pi\left(x, y, \underline{\theta}\right)$ for every (x, y). The target social choice rule satisfies i) $\hat{y}\left(\underline{\theta}\right) > \inf Y$ and ii) type $\overline{\theta}$ agent's (IC) constraint is binding while type $\underline{\theta}$ agent's (IC) constraint is slack. This second condition, together with SCP, implies that $\hat{x}\left(\underline{\theta}\right) < \hat{x}\left(\overline{\theta}\right)$.

The next proposition shows that no per-unit fee schedule achieves intermediated implementation in Example 1:

Proposition 1. *In Example 1, under any per-unit fee schedule, there exists no subgame perfect equilibrium of the contracting game where every agent consumes his target consumption bundle.*

Proposition 1 conveys a disturbing message: invoice-price schedules may not achieve the manufacturer's goal if the distribution cost depends directly on the customer's hidden taste; income taxation may not yield the desired amount of redistribution if firms cannot directly contract on the worker's effective labor; finally, per-unit coverage subsidies may fail to induce private companies to provide the desired amount of insurance.

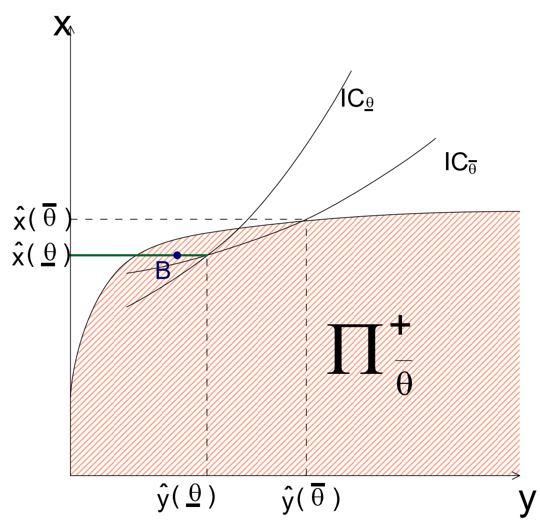

Figure 3: $\hat{\psi}_{per-unit}$ cannot achieve intermediated implementation when intermediaries have interdependent values.

To better understand this negative result, notice first that $\hat{\psi}_{per-unit}$ cannot achieve intermediated implementation in the current setting. As illustrated by Figure 3, now since intermediaries prefer to serve type $\overline{\theta}$ agent than type $\underline{\theta}$ agent, other things being equal, we can no longer separate region $\Pi^{+}_{\overline{\theta}}$ where each bundle makes a profit when it is served to type $\overline{\theta}$ agent, from region U^{+} where each agent's utility is lower bounded by its counterpart under the target social choice rule. To see why this non-separability is problematic, notice that in any equilibrium where each agent consumes his target bundle, any intermediary i can add $(\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}) - \epsilon)$ to its menu and attract both types of agents. When ϵ is small, i's gross profit is given approximately by $\rho\pi\left(\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}), \overline{\theta}\right) + (1-\rho)\pi(\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}), \underline{\theta})$, where ρ stands for the population of type $\overline{\theta}$ agents in the market. In the meantime, i pays a low fee $\pi(\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}), \underline{\theta})$ to the principal and makes a profit overall.[12]

Thus under any per-unit fee schedule, any equilibrium outcome achieves intermediated implementation must involve cross-subsidization, whereby an intermediary incurs a loss from serving the target consumption bundle (denoted by b) to some type of agents, makes a profit from serving the target bundle (denoted by b') to the other type of agents, and breaks even on average. To begin with, consider any equilibrium outcome where $b = (\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}))$ and $b' = \left(\hat{x}\left(\overline{\theta}\right), \hat{y}\left(\overline{\theta}\right)\right)$. Take any intermediary i that serves b to type $\underline{\theta}$ agent and notice the following contradiction: on the one hand, if i is the only intermediary that serves b to type $\underline{\theta}$ agent, then all $j \neq i$ are serving b' to type $\overline{\theta}$ agent and making a profit, and i can offer $\left(\hat{x}\left(\overline{\theta}\right), \hat{y}\left(\overline{\theta}\right) - \epsilon\right)$ to type $\overline{\theta}$ agent and increase its profit; on the other hand, if another intermediary $j \neq i$ is serving b to type $\underline{\theta}$ agent, too, i can drop b from its menu and save the loss. In the appendix, we complete this argument and show that no per-unit fee schedule achieves intermediated implementation in the current example.[13]

3.1.3 Monopoly Power

In general, per-unit fee schedules cannot implement the target social choice rule when the intermediary has monopoly power. The reason is straightforward: under most per-unit fee schedules of interest, the monopolistic intermediary can distort the allocations of some agent types in order to extract information rents from other agent types as

[12] Proposition 1 differs from Rothschild and Stiglitz (1976) in two major aspects. First, we study the implementability of any social choice rule that is incentive compatible, individually rational and feasible among agents, whereas Rothschild and Stiglitz (1976) examines the outcome of a laissez-faire economy. Second, our impossibility result holds even if intermediaries can offer menus of bundles, whereas Rothschild and Stiglitz (1976) limits players to offering one single bundle.

[13] Miyazaki (1977), Wilson (1977), Riley (1979) and Netzer and Scheuer (2014) achieve cross-subsidization under a different solution concept where players can anticipate the response of other players to their contractual offers.

in the case of monopolistic screening.[14] This is best illustrated by Figure 4, where the monopolistic intermediary implements the social choice rule (x^*, y^*) — (part of) which lies in region Π^+ — that maximizes its net profit, taking the per-unit fee schedule $\hat{\psi}_{per-unit}$ as given.

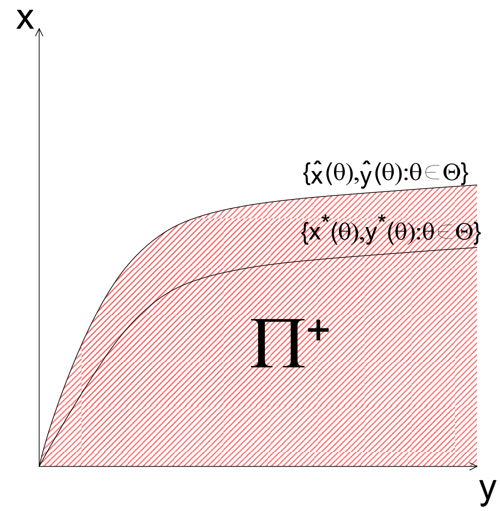

Figure 4: Monopolistic intermediary

3.2 Distribution Regulation

In this section, we show that the distribution regulation $\hat{\psi}_{distr}$ achieves intermediated implementation under mild regularity conditions, and this is true even if intermediaries have interdependent values or monopoly power. Throughout this section, let x denote consumption goods and y denote monetary transfers, and suppose the agent's payoff is quasi-linear in monetary transfers:[15]

Assumption 2. $U(x, y, \theta) = v(x, \theta) - y$.

A *consumption rule* $x : \Theta \to X$ is a deterministic mapping from the type space to the consumption space. It is *implementable* if there exists a *transfer rule* $y : \Theta \to Y$ such that the combined social choice rule $(x, y) : \Theta \to X \times Y$ is incentive compatible among agents.

[14]The concern for monopolistic screening is absent from the existing studies on double marginalization, which examine the phenomenon in which firms at different vertical levels in the supply chain have their respective market powers and apply their own markups in prices.

[15]See the online appendix for the opposite case where x represents monetary transfers. See Corollary 2 for an extension to CARA utility.

3.2.1 Competitive Intermediaries

This section examines the case of where intermediaries are perfectly competitive. The results presented below hold in both cases of private values and interdependent values. Throughout this section, let the agent's type space be finite:

Assumption 3. $|\Theta| = N \in \mathbb{N}$.

In addition, suppose the target consumption rule satisfies a mild regularity condition called (DU). Formally, we say that a bijection $\pi : \Theta \to \Theta$ constitutes a *cyclic permutation over* Θ if $\theta \to \pi(\theta) \to \pi \circ \pi(\theta) \to \cdots$ forms a $|\Theta|$-cycle. A consumption rule satisfies (DU) if the following condition holds:

Definition 3. Permuting a consumption rule $x : \Theta \to X$ yields a distinct total utility among agents (DU) *if for any* $\Theta' \subset \Theta$ *such that* $x(\theta) \neq x(\theta')$ *for all* $\theta, \theta' \in \Theta'$, *and any cyclic permutation* $\pi : \Theta' \to \Theta'$, *we have*

$$\sum_{\theta \in \Theta'} v(x(\theta), \theta) \neq \sum_{\theta \in \Theta'} v(x(\pi(\theta)), \theta). \tag{DU}$$

Assumption 4. $\hat{x} : \Theta \to X$ *satisfies (DU).*

In the remainder of this section, we first state our main result and discuss its implication, then sketch the proof and finally demonstrate how Assumption 4 can be easily satisfied by implementable consumption rules.

Theorem 2. *Suppose that intermediaries are perfectly competitive, that Assumptions 1 - 4 hold, and that agents choose the target consumption bundle and split evenly between intermediaries when* $\sigma_i = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$, $i = 1, \cdots, I$. *Then under* $\hat{\psi}_{distr}$, *the contracting game has a sub-game perfect equilibrium where* $\sigma_i^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$, $i = 1, \cdots, I$.

Theorem 2 shows that under mild regularity conditions, the distribution regulation $\hat{\psi}_{distr}$ implements the target social choice rule through competitive intermediaries in both cases of private values and interdependent values. Thus in reality, regulating the variety of sold products helps achieve the manufacturer's goal whether or not the dealer's distribution cost depends on the customer's hidden taste; setting employment targets for big employers restores the effectiveness of the Mirrleesian tax schedule whether or not firms can contract on the worker's effective labor; finally, Corollary

2 demonstrates how regulating the variety of sold coverage plans induces private companies to provide the desired amount of insurance.

Proof sketch Under $\hat{\psi}_{distr}$, the only kind of deviation that intermediaries can commit is to permute the target consumption goods across agents. Thus if any permuted consumption rule is not implementable and thus cannot arise in any equilibrium outcome (verified in Lemma 1), then each agent must receive his target consumption good in any equilibrium outcome. Base on this result, we then use the competition between intermediaries to show that no intermediary can successfully attract any type of agent without making a loss when all the other intermediaries are charging the target-level prices. This ensures that all target consumption goods are "correctly" priced at the target level.

The next lemma shows that if an implementable consumption rule satisfies (DU), then it is no longer implementable after any cyclic permutation that switches distinct consumption goods across agents.

Lemma 1. *Under Assumptions 2 and 3, if $x : \Theta \to X$ is implementable and satisfies (DU), then for any $\Theta' \subset \Theta$ such that $x(\theta) \neq x(\theta')$ for all $\theta, \theta' \in \Theta'$, and any cyclic permutation π over Θ', $x \circ \pi : \Theta' \to X$ is not implementable among agents whose type belongs to Θ'.*

To better understand Lemma 1, we first revisit Example 1 where X and Θ are single-dimensional and the agent's utility of consumption satisfies SCP in (x, θ). In this setting, Milgrom and Shannon (1994) shows that a consumption rule x is implementable if and only if it is non-decreasing in the agent's type (M). Once we limit attention to those agents who receive distinct consumption goods, then any implementable consumption rule is strictly increasing in the agent's type. In this case, permuting consumption goods across agents destroys (M), yielding a new consumption rule that is no longer implementable.

In general, it is Rochet (1987)'s cyclic monotonicity (CMON) that characterizes the implementability of a consumption rule. To relate Lemma 1 to (CMON), suppose the economy consists of two types of agents θ_1 and θ_2 who obtain different consumption goods x_1 and x_2 under the target social choice rule. Since the target consumption rule is implementable and thus satisfies (CMON), it follows that $v(x_1, \theta_1)+v(x_2, \theta_2) \geq v(x_1, \theta_2) + v(x_2, \theta_1)$. Meanwhile, if permuting the target consumption goods across agents yields a new implementable rule, then $v(x_2, \theta_1)+v(x_1, \theta_2) \geq v(x_1, \theta_2)+v(x_2, \theta_1)$.

Thus the permuted consumption rule is not implementable if the target consumption rule satisfies (DU), i.e., $v(x_1, \theta_1) + v(x_2, \theta_2) \neq v(x_1, \theta_2) + v(x_2, \theta_1)$.

When does (DU) hold? Examples below show that Assumption 4 can be easily satisfied by implementable consumption rules.

Example 2. *If X and Θ are totally-ordered sets and $v(x, \theta)$ is supermodular in x and θ, then any implementable consumption rule satisfies (M) and hence (DU).*

Example 3. *Let $\Theta, X \subset \mathbb{R}^d$ and write $\Theta = \left\{\vec{\theta}_1, \cdots, \vec{\theta}_N\right\}$. Suppose for illustrative purpose only that $v\left(\vec{x}, \vec{\theta}\right) = \vec{\theta} \cdot \vec{x}$. In this case, a consumption rule is an $N \times d$-dimensional vector $(\vec{x}_1, \cdots, \vec{x}_N)$ where $\vec{x}_i$ denotes type $\vec{\theta}_i$ agent's consumption good. By definition, a target consumption rule satisfies (DU) if*

- *First, it does not belong to any $(2 \times d - 1)$-dimensional subspace $\vec{\theta}_i \cdot \vec{x}_i + \vec{\theta}_j \cdot \vec{x}_j = \vec{\theta}_i \cdot \vec{x}_j + \vec{\theta}_j \cdot \vec{x}_i$, $i, j \in \{1, \cdots, N\}$.*
- *Second, it does not belong to any $(3 \times d - 1)$-dimensional subspace $\vec{\theta}_i \cdot \vec{x}_i + \vec{\theta}_j \cdot \vec{x}_j + \vec{\theta}_k \cdot \vec{x}_k = \vec{\theta}_{\pi(1)} \cdot \vec{x}_{\pi(2)} + \vec{\theta}_{\pi(2)} \cdot \vec{x}_{\pi(3)} + \vec{\theta}_{\pi(3)} \cdot \vec{x}_{\pi(1)}$ where $i, j, k \in \{1, \cdots, N\}$ and π is any bijection between $\{1, 2, 3\}$ and $\left\{\vec{\theta}_i, \vec{\theta}_j, \vec{\theta}_k\right\}$.*
- $\cdots$
- *Finally, it does not belong to any $(N \times d - 1)$-dimensional subspace $\vec{\theta}_1 \cdot \vec{x}_1 + \cdots \vec{\theta}_N \cdot \vec{x}_N = \vec{\theta}_{\pi(1)} \cdot \vec{x}_{\pi(2)} + \cdots \vec{\theta}_{\pi(N)} \cdot \vec{x}_{\pi(1)}$ where π is any bijection between $\{1, 2, \cdots, N\}$ and Θ.*

When X^N is a rich subset of $\mathbb{R}^{N \times d}$, these conditions impose virtually no restriction on the set of implementable consumption rules that belong to $X^{N \times d}$.

CARA utility The next corollary extends Lemma 1 and Theorem 2 to CARA Bernoulli utility functions and demonstrates the applicability of the results so far to decentralized insurance sales.

Corollary 2. *Lemma 1 and Theorem 2 hold in the setting of Section 2.2.3, if agents have a CARA Bernoulli utility function $u(c) = -\frac{1}{\lambda} \exp(-\lambda c)$.*

Coarse distribution regulation In reality, the principal may not be able to enforce the exact distribution requirement for practical reasons such as the lack of precise knowledge about agent type distribution, complexity concerns and administrative barriers. In Section 5, we detail these reasons and demonstrate how coarse distribution may help solve the problems they create.

Intermediary's technology The maintained assumption so far is that intermediaries are homogeneous and the technology has constant returns to scale. In Appendix B.3, we extend the analysis to heterogeneous intermediaries and non-constant-returns-to-scale technologies.

3.2.2 Monopoly Power

This section examines the case of monopolistic intermediary. The result below holds in both cases of private values and interdependent values. Throughout this section, we make a few standard assumptions in the mechanism design literature:

Assumption 5. *$X \subset \mathbb{R}$, $\Theta = \left[\underline{\theta}, \overline{\theta}\right] \subset \mathbb{R}$ and P_θ has a positive probability density function p_θ.*

Assumption 6. *$v(x,\theta)$ has SCP in (x,θ) and is continuously differentiable in θ for every x; $v_\theta(x,\theta)$ is bounded uniformly across (x,θ).*

Define $\hat{P}_y$ and μ_y in the same way as we did with $\hat{P}_x$ and μ_x.[16] For any $q \in [0,1]$, let $q(\mu)$ denote the $100q^{th}$ percentile of any probability distribution μ over $\mathbb{R}$. Fix an arbitrary $q \in [0,1]$ and consider a policy that charges $\hat{\psi}_{per-unit}$ if first, the intermediary matches the distribution over consumption goods under the target social choice rule; and second, the $100q^{th}$ percentile of the charge prices is upper bounded by its counterpart under the target social choice rule:[17]

$$\breve{\psi}_{distr}\left(\mu_x, q\left(\mu_y\right)\right) = \begin{cases} \hat{\psi}_{per-unit}\left(\mu_x\right) & \text{if } \mu_x = \hat{P}_x \text{ and } q\left(\mu_y\right) \leq q\left(\hat{P}_y\right), \\ +\infty & \text{otherwise.} \end{cases} \tag{3.3}$$

Theorem 3. *Suppose that the intermediary is monopolistic, and that Assumptions 1, 2, 5 and 6 hold. Under $\breve{\psi}_{distr}$, we have $\sigma^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$.*

[16]Recall that $\hat{P}_x$ denotes the probability measure over $(X, \mathcal{X})$ under the target social choice rule, and μ_x denotes the measure over $(X, \mathcal{X})$ that is induced by the intermediary's sold bundles.

[17]See Appendix B.2 for an extension to coarse distribution regulation.

Theorem 3 shows that regulating the variety of sold consumption goods, together with a price cap, induces the monopolistic intermediary to offer the full menu of target consumption bundles. Besides the aforementioned applications, this result explains why the housing authorities in the U.S. impose rigid rulings on major real-estate companies regarding the variety of housing types that can be developed for low-income households. In addition, it suggests that one way to stop schools from oversubscribing students into the National School Lunch program is to limit the percentage of eligible students based on the income and demographic distribution within the school district.[18]

To better understand Theorem 3, notice that under Assumptions 5 and 6, any implementable consumption rule is non-decreasing in the agent's type. This result, together with the requirement that $\mu_x = \hat{P}_x$, limits the principal to offering $x^* = \hat{x}$. Now from the envelope theorem (Milgrom and Segal (2002)), it follows that

$$y^*(\theta) = v\left(\hat{x}(\theta), \theta\right) - \int_{\underline{\theta}}^{\theta} v_\theta\left(\hat{x}(s), s\right) ds + \text{constant}_1.$$

In addition, since

$$\hat{y}(\theta) = v\left(\hat{x}(\theta), \theta\right) - \int_{\underline{\theta}}^{\theta} v_\theta\left(\hat{x}(s), s\right) ds + \text{constant}_2,$$

it follows that the intermediary — whose profit is increasing in y — will put $y^* = \hat{y}$ in order to satisfy $q\left(\mu_y\right) \leq q\left(\hat{P}_y\right)$.

Discrete types When agent types are discrete, we cannot use the envelope theorem to pin down the transfer rule as a function of the consumption rule. In these situations, the effectiveness of $\breve{\psi}_{distr}$ depends on which of the agent's incentive constraints are binding under the target social choice rule. To illustrate, suppose $\Theta = \left\{\overline{\theta}, \underline{\theta}\right\}$ and consider $\breve{\psi}_{distr}$ with $q = 0$. A close inspection of Figures 5 and 6 reveals two patterns. First, our policy implements the target social choice rule if $\text{IC}_{\overline{\theta}}$ is binding. Second, if $\text{IC}_{\overline{\theta}}$ is slack, then the intermediary will offer $\left(\hat{x}\left(\overline{\theta}\right), \overline{y}\left(\overline{\theta}\right)\right)$ to type $\overline{\theta}$ agent instead; the resulting distortion, measured by the difference between $\overline{y}\left(\overline{\theta}\right)$ and $\hat{y}\left(\overline{\theta}\right)$, depends on a number of factors such as the distance between the target consumption bundles

[18]National School Lunch program is a federal subsidized lunch program for students with private needs. For further references, see David Bass, "Fraud in the Lunchroom?," *Education Next*, Winter 2010, Volume 10, Issue 1.

and the elasticity of the agents' indifference curves.

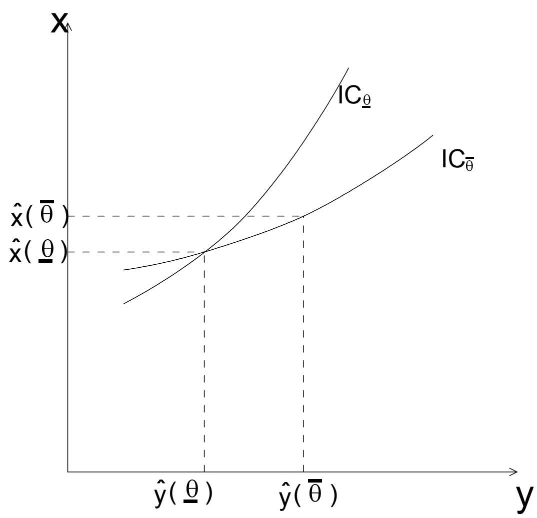

Figure 5: $IC_{\overline{\theta}}$ is binding.

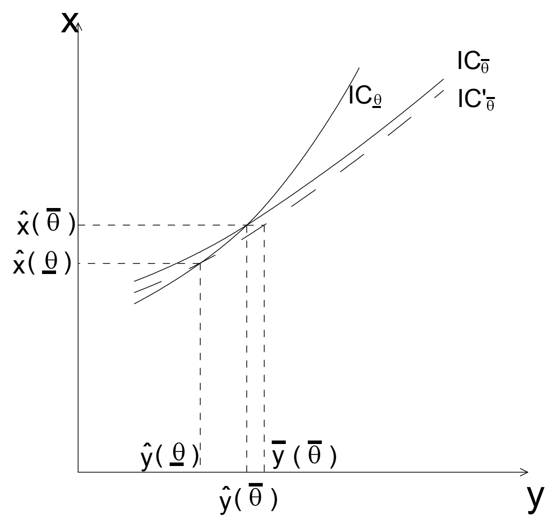

Figure 6: $IC_{\overline{\theta}}$ is slack.

4 Extensions

Turning interdependent-value environments into private-value environments We propose two organizational solutions that turn interdependent-value environments into private-value environments. First, notice that intermediaries have private values if parties can contract directly on the variables that depend on the agent's type. In the example of income taxation, this means that the government should help firms improve the monitoring technology in order to compensate workers based on effective labor rather than working hours.

Second, we have so far allowed the principal's objective function to depend directly on the agent's type. This observation, albeit a simple one, suggests a second way to eliminate interdependent values: let the principal internalize the part of the intermediary's payoff that depends directly on the agent's type. In reality, this means that single-payer health care systems do not face the same challenge that interdependent values create as market-based systems do. And manufacturers should handle maintenance and repair directly if the cost of providing these services varies much with the customer's level of savviness.

Knowledge requirement and coarse distribution regulation In order to implement $\hat{\psi}_{per-unit}$, the principal needs to know the target social choice rule and the intermediary's payoff under the target social choice rule. In order to enforce $\hat{\psi}_{distr}$, she also needs to know the distribution $\hat{P}_x$ over x under the target social choice rule.

In reality, the principal may not have perfect knowledge about the agent's type distribution P_θ. This limitation has two effects on the outcome of intermediated implementation. First, the principal may revise her target to those social choice rules that have a coarse distribution over x. Second, for any given target social choice rule $\hat{x}$, the principal may not be able to compute the exact distribution requirement $\hat{P}_x$. The analysis below demonstrates how coarse distribution requirements may help solve the second problem.

Formally, take as given any social choice rule $(\hat{x}, \hat{y})$ that satisfies (IC) and (IR) and is feasible. Suppose the principal knows $\hat{\pi}(x)$ for every $x \in \hat{x}(\Theta)$, but can only enforce $\left\| \frac{\mu_{i,x}}{\int_{x \in \hat{x}(\Theta)} d\mu_{i,x}} - \hat{P}_x \right\| < \epsilon$ for some small $\epsilon > 0$ ($\| \cdot \|$ denotes the supremum norm). Consider the following policy, which charges a fee $\hat{\psi}_{per-unit}$ if the probability distribution over sold x is ϵ-close to $\hat{P}_x$, and inflicts a big penalty on the intermediary otherwise:

$$\tilde{\psi}_{distr}\left(\mu_{i,x}\right) = \begin{cases} \hat{\psi}_{per-unit}(\mu_{i,x}) & \text{if } \left\| \frac{\mu_{i,x}}{\int_{x \in \hat{x}(\Theta)} d\mu_{i,x}} - \hat{P}_x \right\| < \epsilon, \\ +\infty & \text{otherwise.} \end{cases} \tag{4.1}$$

The next corollary shows that Theorem 2 remains true when ϵ is sufficiently small:

Corollary 3. *Suppose that intermediaries are competitive, that Assumptions 1 - 4 hold, that agents choose their target bundle in case of indifference, and that they split evenly between intermediaries who offer the same bundle. When $\epsilon > 0$ is sufficiently small, the contracting game induced by $\tilde{\psi}_{distr}$ has a sub-game perfect equilibrium where $\sigma_i^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$, $i = 1, \cdots, I$.*

In Appendix B.2, we show that under mild regularity conditions, coarse distribution regulations enables the implementation of slightly distorted consumption rules in the case of monopolistic intermediary.

5 Conclusion

Motivated by real-world problems such as sales, taxation and health care regulation, we propose a framework for studying a new class of implementation problems called intermediated implementation. In these problems, one or more intermediaries can specify the full menu of the multi-faceted bundles that they offer to agents, whereas the principal is limited to regulating some but not all aspects of the consumed bundles, due to legal, information and administrative barriers. Our main research question concerns how the principal can implement through intermediaries the target social choice rule that maximizes her objective function, subject to agents' incentive compatibility constraint, individual rationality constraint and feasibility constraint.

The two policies of our interest: per-unit fee schedule and distribution regulation, have natural counterparts in reality. The main result shows that the effectiveness of these policies depends on whether intermediaries have private values or interdependent values, and whether the intermediary market is competitive or monopolistic. Our proof strategy leverages different characterizations of incentive compatible social choice rules that range from the rationality of players to monotonicity and cyclic monotonicity. Applications to sales, taxation and health care regulation are discussed. Various extensions to heterogeneous intermediaries, non-constant returns to scale and coarse distribution regulation are finally considered.

In the future, it would be interesting to consider other kinds of industrial organizations (e.g., horizontal differentiation), to introduce market frictions to the analysis (e.g., search friction), to characterize the second-best per-unit fee schedule in the case of interdependent values or monopoly power, and perhaps to characterize all policies that achieve intermediated implementation.

A Mathematical Proofs

Proof of Theorem 1

Proof. When $\sigma_j^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ $\forall j \neq i$, suppose there exists $(\hat{x}(\theta'), y) \in \sigma_i^*$ where $u(\hat{x}(\theta'), y, \theta) \geq u(\hat{x}(\theta), \hat{y}(\theta), \theta)$ for some $\theta \in \Theta$, $\pi(\hat{x}(\theta'), y) \geq \psi(\hat{x}(\theta'))$, and one of these inequalities is strict. W.l.g.o., suppose that $u(\hat{x}(\theta'), y, \theta) > u(\hat{x}(\theta), \hat{y}(\theta), \theta)$. Since $(\hat{x}, \hat{y})$ is an incentive compatible social choice rule, it follows that $u(\hat{x}(\theta), \hat{y}(\theta), \theta) \geq$

$u\left(\hat{x}\left(\theta'\right),\hat{y}\left(\theta'\right),\theta\right)$ and hence that $u\left(\hat{x}\left(\theta'\right),y,\theta\right) > u\left(\hat{x}\left(\theta'\right),\hat{y}\left(\theta'\right),\theta\right)$. Under Assumption 1, it follows that $y < \hat{y}\left(\theta'\right)$ and hence that $\psi\left(\hat{x}\left(\theta'\right)\right) \leq \pi\left(\hat{x}\left(\theta'\right),y\right) < \pi\left(\hat{x}\left(\theta'\right),\hat{y}\left(\theta'\right)\right) = \psi\left(\hat{x}\left(\theta'\right)\right)$, a contradiction. $\square$

Proof of Corollary 1

Proof. Take any equilibrium of the contracting game. Let A be the set of the bundles that agents end up consuming, and $A_i \subset A$ be the bundles that are served by intermediary i. For any $S \subset A$, use $\mu(S)$ to denote the measure of agents who end up consuming the bundles in S.

We proceed in two steps. First, suppose, to the contrary, that a positive measure of agents consume bundles that allow intermediaries to make a profit, i.e., $\sup_{(x,y)\in A} \pi(x,y) \triangleq \overline{\pi} > 0$. Take any $\underline{\pi} \in (0,\overline{\pi})$. Let $B = \{(x,y) \in A : \pi(x,y) \in [\underline{\pi},\overline{\pi}]\}$. Consider two cases:

1. $\mu(B\backslash A_i) > 0$ for some i. In this case, i can add $B\backslash A_i$ to its menu and increase its profit.

2. $\mu(B\backslash A_i) = 0$ for all $i = 1,\cdots,I$. In this case, suppose intermediary i undercuts the y-dimension of every bundle in B by a small amount ϵ. This deviation has two effects: first, it helps i attract all agents who used to purchase from $B \cap A_{-i}$, resulting in a gain that is bounded below by $\mu(B \cap A_{-i}) \cdot \underline{\pi}$; second, it undercuts i's own customers, which leads to a loss of $\mathcal{O}(\epsilon)$ under the assumption of private values. The net change in i's payoff is positive when ϵ is small.

Thus in every equilibrium, any positive measure of sold bundles yields zero profit.

Second, suppose there exists a positive measure of agents in $\Theta' \subset \Theta$ whose equilibrium utility falls short of its target-level counterpart. Then any intermediary i can add $\{(\hat{x}(\theta),\hat{y}(\theta)+\epsilon) : \theta \in \Theta'\}$ where ϵ is positive but small to its menu and make a profit. Thus, every agent obtains his target level of utility in every equilibrium. $\square$

Proof of Proposition 1

Proof. In Section 3.1.2, we have shown that it is impossible to achieve intermediated implementation through $\hat{\psi}_{per-unit}$ or any $\psi_{per-unit}$ under which $b = \left(\hat{x}\left(\underline{\theta}\right),\hat{y}\left(\underline{\theta}\right)\right)$ and $b' = \left(\hat{x}\left(\overline{\theta}\right),\hat{y}\left(\overline{\theta}\right)\right)$ (recall that b (resp. b') denotes the target consumption bundle that incurs a loss (resp. makes a profit) when it is being served to its target agent).

Now suppose intermediated implementation can be achieved through $\psi_{per-unit}$, under which $b = \left(\hat{x}\left(\overline{\theta}\right), \hat{y}\left(\overline{\theta}\right)\right)$ and $b' = \left(\hat{x}\left(\underline{\theta}\right), \hat{y}\left(\underline{\theta}\right)\right)$. Take any intermediary i who offers b to type $\overline{\theta}$ agent and notice the following contradiction: on the one hand, if i is the only intermediary that serves b to type $\overline{\theta}$ agent, then all $j \neq i$ are serving b' to type $\underline{\theta}$ and making a profit, and i can offer $\left(\hat{x}\left(\overline{\theta}\right), \hat{y}\left(\overline{\theta}\right) - \epsilon\right)$ to both types of agents and increase its profit; on the other hand, if some $j \neq i$ is offering b to type $\overline{\theta}$ agent, too, i can drop b from its menu and save a loss. □

Proof of Lemma 1

Proof. Take any implementable consumption rule $x : \Theta \to X$ and any $\Theta' \subset \Theta$ such that $x(\theta) \neq x(\theta')$ for all $\theta, \theta' \in \Theta'$. Denote $|\Theta'| = m$. Write any cyclic permutation over Θ' as $\pi(1) \to \pi(2) \to \cdots \to \pi(m) \to \pi(1)$, where π is any arbitrary bijection between $\{1, 2, \cdots, m\}$ and Θ'. For convenience, write $\pi(i) = \theta_i$, $i = 1, \cdots, m$.

First, since $x : \Theta' \to X$ is implementable among agents in Θ', there exists a transfer rule $y : \Theta' \to \mathbb{R}$ such that

$$\begin{aligned}
v(x_2, \theta_1) - y_2 &\leq v(x_1, \theta_1) - y_1, \\
v(x_3, \theta_2) - y_3 &\leq v(x_2, \theta_2) - y_2, \\
&\cdots \\
v(x_1, \theta_m) - y_1 &\leq v(x_m, \theta_m) - y_m.
\end{aligned}$$

Summing over these inequalities yields $\sum_{i=1}^{m} v(x_{i+1}, \theta_i) \leq \sum_{\theta \in \Theta'} v(x(\theta), \theta)$.

Second, if $(x_2, \cdots, x_m, x_1)$ is implementable among $(\theta_1, \cdots, \theta_m)$, then there exists a transfer rule $y' : \Theta' \to \mathbb{R}$ such that

$$\begin{aligned}
v(x_2, \theta_1) - y'_2 &\geq v(x_1, \theta_1) - y'_1, \\
v(x_3, \theta_2) - y'_3 &\geq v(x_2, \theta_2) - y'_2, \\
&\cdots \\
v(x_1, \theta_m) - y'_1 &\geq v(x_m, \theta_m) - y'_m.
\end{aligned}$$

Summing over these inequalities yields $\sum_{i=1}^{m} v(x_{i+1}, \theta_i) \geq \sum_{\theta \in \Theta'} v(x(\theta), \theta)$. Thus, in order for $(x_2, \cdots, x_m, x_1)$ to be not implementable among $(\theta_1, \cdots, \theta_m)$, it suffices to have $\sum_{i=1}^{m} v(x_{i+1}, \theta_i) \neq \sum_{\theta \in \Theta'} v(x(\theta), \theta)$. Repeating this argument for all $\Theta' \subset \Theta$ and all cyclic permutations over Θ' yields the result. □

Proof of Theorem 2

Proof. Let $x^i : \Theta \to X$ be the consumption correspondence that is induced by the bundles that intermediary i sells, and use $x^i(\Theta)$ to denote the image of Θ under x^i. Apparently, any x^i that meets the distribution requirement must satisfy $x^i(\Theta) = \hat{x}(\Theta)$. For any $x, x' \in \hat{x}(\Theta)$ such that $\hat{x}(\theta) = x$ and $x' \in x^i(\theta)$ for some θ and i, we say that x (resp. x') is an immediate predecessor (resp. successor) of x' (resp. x), and write $x \to x'$.

Fix $\sigma_j^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ $\forall j \neq i$, under which i can meet the distribution requirement by offering σ_i^*. In the remainder of this proof, let $y^j(x)$ denote the price that intermediary j charges for $x \in \hat{x}(\Theta)$. By assumption, we have $y^j(x) = \hat{y}(x)$ for all $j \neq i$ (recall that $\hat{y}(x)$ denotes the unique $y \in Y$ such that $(x, y) \in (\hat{x}, \hat{y})(\Theta)$).

Take any σ_i and consider the graph that is generated by (σ_i, σ_j^*). For every $x \in \hat{x}(\Theta)$, only three situations can happen: (a) x has no immediate predecessor or successor, (b) x is part of a cycle, and (c) x is part of a chain. We now go through these cases one by one.

Case (a) In this case, each agent obtains his target consumption good. From this follows that $y^i(x) = \hat{y}(x)$, as the remaining cases can be ruled out as follows:

- If $y^i(x) > \hat{y}(x)$, then all agents whose target consumption good equals x will purchase x from $j \neq i$. Thus i cannot satisfy the distribution requirement and hence is not best-responding to σ_{-i}^*.

- If $y^i(x) < \hat{y}(x)$, then all agents whose target consumption good equals x will purchase x from i. Then in order to satisfy the distribution requirement, i must charge $y^i(x') \leq \hat{y}(x')$ for each $x' \in \hat{x}(\Theta)$ and hence incur a loss, a contradiction.

Case (b) This case is ruled out by (DU). For any cycle $x \to x' \to \cdots$, let $\theta, \theta', \cdots$ be any sequence of agents who consume (1) x, x', $\cdots$ under the target social choice rule and (2) x', x'', $\cdots$ in any equilibrium outcome. Since $\hat{x} : \Theta \to X$ satisfies (DU), the new consumption rule that assigns x' to θ, x'' to θ', so on and so forth, cannot be part of any incentive compatible allocation and hence cannot arise in any equilibrium outcome, a contradiction.

Case (c) This case is impossible, too. Specifically, take any finite chain with an end node $x' = \hat{x}(\theta')$ and let $x'' = \hat{x}(\theta'')$ be an immediate predecessor of x'. Notice three things:

1. Incentive compatibility means that $v(x'', \theta'') - \hat{y}(x'') \geq v(x', \theta'') - \hat{y}(x')$;

2. $x'' \to x'$ means that $v(x'', \theta'') - \hat{y}(x'') \leq v(x', \theta'') - y^i(x')$;

3. The very definition of end node means that all type θ' agents consume x'.

From these observations, it follows that $y^i(x') = \hat{y}(x')$, as the remaining cases can be ruled out as follows:

- If $y^i(x') > \hat{y}(x')$, then $v(x', \theta'') - y^i(x') < v(x', \theta'') - \hat{y}(x') \leq v(x'', \theta'') - \hat{y}(x'')$ and hence type θ'' agent will not choose $(x', y^i(x'))$, a contradiction.

- If $y^i(x') < \hat{y}(x')$, then the fact that x' is an end node means that all type θ' agents purchase x' from intermediary i. This, together with $x'' \to x'$, implies that i is violating the distribution requirement and hence is not best-responding to σ^*_{-i}, a contradiction.

Repeating this argument for every x along the chain yields that $y^i(x) = \hat{y}(x)$ for all x. Thus $\sigma_i = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\} = \sigma^*_j$ and $x'' \to x'$ cannot happen under our tie-breaking rule. □

Proof of Corollary 2

Proof. In the setup of Section 2.2.3, if agents have a CARA Bernoulli utility function $u(c) = -\frac{1}{\lambda}\exp(-\lambda c)$, then (IC) can be rewritten as $v(x(\theta), \theta) - \lambda y(\theta) \leq v(x(\theta'), \theta) - \lambda y(\theta')$, where $v(x, \theta) \triangleq \log\left(-\sum_{s=1}^{d} \theta_s \exp(-\lambda x_s)\right)$. Plugging this into the proofs of Lemma 1 and Theorem 2 yields the result. □

Proof of Corollary 3

Proof. Take $\epsilon \ll \min_{\theta \in \Theta} P_\theta(\theta)$ and fix $\sigma^*_j = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ for all $j \neq i$. In the proof of Theorem 2, we show that every $x \in \hat{x}(\Theta)$ is either (a) separated from the remaining graph, (b) part of a cycle, or (c) part of a chain. In the current setting, it is easy to verify that $\sigma_i = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ in Case (a), that Cases (b) is impossible because of the assumption of (DU), and that Case (c) is impossible under our tie-breaking rule. □

Proof of Theorem 3

Proof. Follows closely the proof sketch in Section 3.2.2 and hence is omitted. □

B Online Appendix

B.1 Distribution Regulation: Transfers

In the setup of Sections 3.2.1 and 3.2.2, now let x denote monetary transfers and y denote consumption goods.

Assumption 7. *$U(x, y, \theta) = x - v(y, \theta)$, where $v(y, \theta)$ satisfies SCP and is continuously differentiable in y for every θ, and $v_\theta(y, \theta)$ is uniformly bounded across (y, θ).*

Assumption 8. *The ordinary differential equation $v_y(y(\theta), \theta)y'(\theta) = \frac{d}{d\theta}\hat{x}(\theta)$ with an initial condition $y\left(q\left(P_\theta\right)\right) = \hat{y}\left(q\left(P_\theta\right)\right)$ for some $q \in [0, 1]$ has a unique solution $\hat{y}(\cdot)$ on Θ.*

Corollary 4. *Under Assumptions 1, 5 and 7,*

(i) Suppose intermediaries are perfectly competitive. Under $\hat{\psi}_{distr}$, the contracting game has a sub-game perfect Nash equilibrium where $\sigma_i^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$, $i = 1, \cdots, I$;

(ii) Suppose the intermediary has monopoly power and Assumption 8 holds. Under a modified $\breve{\psi}_{distr}$ where the distribution regulation on y is replaced with $q(\mu_y) = q\left(\hat{P}_y\right)$, $\sigma^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$.

Proof. Take any incentive compatible social choice rule (x, y). On both sides of the following first-order incentive constraint:

$$x(\theta) - v(y(\theta), \theta) = -\int_{\underline{\theta}}^{\theta} v_\theta(y(s), s)ds + \text{constant},$$

taking total derivative with respect to θ yields

$$x'(\theta) = v_y(y(\theta), \theta)y'(\theta).$$

Since $y(\theta)$ is non-decreasing in θ (Milgrom and Shannon (1994)) and $v(y,\theta)$ is increasing in y (Assumption 1), it follows that $x(\theta)$ is non-decreasing in θ. Thus the only transfer rule that satisfies the distribution requirement is $\hat{x}$.

Part (i): the above discussion suggests that i cannot make a profit and satisfy the distribution requirement simultaneously given σ^*_{-i}, and hence that σ^*_i is a best-response to σ^*_{-i}.

Part (ii): the social choice rule $(\hat{x}, y)$ proposed by the monopolistic intermediary satisfies

$$\frac{d}{d\theta}\hat{x}(\theta) = v_y\left(y(\theta),\theta\right)\frac{d}{d\theta}y(\theta),$$

and

$$y\left(q(P_\theta)\right) = q(\mu_y) = q\left(\hat{P}_y\right) = \hat{y}\left(q(P_\theta)\right).$$

Under Assumption 8, it follows that $y = \hat{y}$. $\square$

B.2 Coarse Distribution Regulation Against Monopolistic Intermediary

In the setup of Section 3.2.2, suppose that the c.d.f. of agent types (denoted by F) has a p.d.f. that is bounded uniformly below by a constant $f > 0$, and that $\hat{x}$ is strictly increasing and is c-Lipschitz for some $c > 0$. For every $x \in X$, define $\hat{P}_x(x) = F\left(\theta : \hat{x}(\theta) \le x\right)$, $P^*_x(x) = F\left(\theta : x^*(\theta) \le x\right)$ and $P^*_{x,-}(x) = F\left(\theta : x^*(\theta) < x\right)$. Consider the following distribution regulations:

$$\left\|P^*_x - \hat{P}_x\right\|, \left\|P^*_{x,-} - \hat{P}_x\right\| < \varepsilon, \tag{B.1}$$

where $\|\cdot\|$ denotes the sup-norm.

Corollary 5. *Suppose that the intermediary is monopolistic, that Assumptions 1, 2, 5 and 6 hold, and that the above described assumptions are satisfied. Then* $\|x^* - \hat{x}\| < \frac{c\varepsilon}{f}$ *under* (B.1).

Proof. Since $\hat{x}$ is increasing and x^* is non-decreasing, it follows that

$$\varepsilon > \left|\hat{P}_x(x) - P^*_x(x)\right| = \left|F\left(\hat{x}^{-1}(x)\right) - F\left(\sup\left\{\theta : x^*(\theta) = x\right\}\right)\right|,$$

and

$$\varepsilon > \left|\hat{P}_x(x) - P^*_{x,-}(x)\right| = \left|F\left(\hat{x}^{-1}(x)\right) - F\left(\inf\{\theta : x^*(\theta) = x\}\right)\right|.$$

Telescoping yields

$$\begin{aligned} F\left(\hat{x}^{-1}(x)\right) - \varepsilon &\leq F\left(\inf\{\theta : x^*(\theta) = x\}\right) \\ &\leq F\left(\sup\{\theta : x^*(\theta) = x\}\right) \leq F\left(\hat{x}^{-1}(x)\right) + \varepsilon. \end{aligned}$$

Under the aforementioned assumptions, applying $\hat{x} \circ F^{-1}$ to this chain of inequalities yields

$$x - \frac{c\varepsilon}{f} < \hat{x}\left(\inf\{\theta : x^*(\theta) = x\}\right) \leq \hat{x}\left(\sup\{\theta : x^*(\theta) = x\}\right) < x + \frac{c\varepsilon}{f}.$$

Thus any type θ agent who purchases $x^*(\theta) = x$ from the monopolistic intermediary obtains $\hat{x}(\theta) \in \left(x - \frac{c\varepsilon}{f}, x + \frac{c\varepsilon}{f}\right)$ under the target social choice rule. $\square$

B.3 Heterogeneous intermediaries and Non-Constant Returns To Scale

In the setup of Section 3.2.1, let $u(x, y, \theta)$ be agent θ's utility from consuming a bundle (x, y) as before and let $I \geq 2$. In the case where the bundles sold by intermediary i induces a measure μ_i over $(X \times Y \times \Theta, \mathcal{X} \otimes \mathcal{Y} \otimes \Sigma)$, intermediary i's profit is given by $\pi_i(\mu_i)$, where $\pi_i(\cdot)$ can differ across $i = 1, \cdots, I$ and exhibit non-constant returns to scale.

Target allocation A target allocation consists of an incentive compatible social choice rule $(\hat{x}, \hat{y}) : \Theta \to X \times Y$ and a profile $\{\hat{\mu}_i\}_{i=1}^I$, where $\hat{\mu}_i\left(\hat{x}(\theta), \hat{y}(\theta), \theta\right)$ prescribes the measure of type θ agents that intermediary i serves $(\hat{x}(\theta), \hat{y}(\theta))$ with, and $\hat{\mu}_i(x, y, \theta) = 0$ if $(x, y) \neq (\hat{x}(\theta), \hat{y}(\theta))$ for all θ. Suppose each type of agent is served by multiple intermediaries, i.e., for each $\theta \in \Theta$, $\hat{\mu}_i\left(\hat{x}(\theta), \hat{y}(\theta), \theta\right), \hat{\mu}_j\left(\hat{x}(\theta), \hat{y}(\theta), \theta\right) > 0$ for some $i \neq j$.

Policy The principal only observes $\mu_{i,x}$, the measure over $(X, \mathcal{X})$ that is induced by i's sold bundles. Consider the following policy, which taxes away i's profit under the target allocation if the measure over x induced by i's sold bundles matches its counterpart under the target allocation, and inflicts a big penalty on i otherwise:

$$\psi_i\left(\mu_{i,x}\right) = \begin{cases} \pi_i\left(\hat{\mu}_i\right) & \text{if } \mu_{i,x} = \hat{\mu}_{i,x}, \\ +\infty & \text{otherwise.} \end{cases} \tag{B.2}$$

Corollary 6. *Suppose Assumptions 2 - 4 hold, and $\hat{\mu}_i\left(\hat{x}(\theta), \hat{y}(\theta), \theta\right)$ type θ agents purchase $\left(\hat{x}(\theta), \hat{y}(\theta)\right)$ from intermediary i when $\sigma_i = \left\{\left(\hat{x}(\theta), \hat{y}(\theta)\right) : \hat{\mu}_i\left(\hat{x}(\theta), \hat{y}(\theta)\right) > 0\right\}$, $i = 1, \cdots, I$. Then under $\{\psi_i\}_{i=1}^I$, there exists a sub-game perfect equilibrium of the contracting game where $\sigma_i^* = \left\{\left(\hat{x}(\theta), \hat{y}(\theta)\right) : \hat{\mu}_i\left(\hat{x}(\theta), \hat{y}(\theta)\right) > 0\right\}$, $i = 1, \cdots, I$.*

Proof. Follows the proof of Theorem 2 step by step. □

References

Bisin, A., and P. Gottardi (2006): Efficient Competitive Equilibria with Adverse Selection, *Journal of Political Economy*, 114, 485-516.

Epstein, L., and M. Peters (1999): A Revelation Principle for Competing Mechanisms, *Journal of Economic Theory*, 88, 119-160.

Guerrieri, V., R. Shimer, and R. Wright (2010): Adverse Selection in Competitive Search Equilibrium, *Econometrica*, 78(6), 1823-1862.

Jackson, M. O., and H. Sonnenschein (2007): Overcoming Incentive Constraints by Linking Decisions, *Econometrica*, 75(1), 241-257.

Martimort, D., and Lars Stole (2002): The Revelation and Delegation Principles in Common Agency Games, *Econometrica*, 70(4), 1659-1673.

Milgrom, P., and I. Segal (2002): Envelope Theorems for Arbitrary Choice Sets, *Econometrica*, 70(2), 583-601.

Milgrom, P., and C. Shannon (1994): Monotone Comparative Statics, *Econometrica*, 62(1), 157-180.

Mirrlees, J. (1971): An Exploration in the Theory of Optimum Income Taxation, *Review of Economic Studies*, 38(2), 175-208.

MIYAZAKI, H. (1977): The Rat Race and Internal Labor Markets, *Bell Journal of Economics*, 8(2), 394-418.

MUSSA, M., AND S. ROSEN (1978): Monopoly and Product Quality, *Journal of Economic Theory*, 18(2), 301-317.

NETZER, N., AND F. SCHEUER (2014): A Game Theoretic Foundation of Competitive Equilibria with Adverse Selection, *International Economic Review*, 55, 399-422.

PRESCOTT, E., AND R. TOWNSEND (1984): Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard, *Econometrica*, 52(1), 21-46.

RAHMAN, D. (2011): Detecting Profitable Deviations, *Working Paper*, University of Minnesota.

RILEY, J. (1979): Informational Equilibrium, *Econometrica*, 47(2), 331-359.

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

ROTHSCHILD, R., AND J. E. STIGLITZ (1976): Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information, *Quarterly Journal of Economics*, 90(4), 629-649.

SEGAL, I. (1999): Contracting with Externalities, *Quarterly Journal of Economics*, 114(2), 337-388.

—— (2006): Lecture Notes in Contract Theory, *Unpublished Manuscript*, Stanford University.

STANTCHEVA, S. (2014): Optimal Income Taxation with Adverse Selection in the Labor Market, *Review of Economic Studies*, 81, 1296-1329.

WILSON, C. (1977): A Model of Insurance Markets with Incomplete Information, *Journal of Economic Theory*, 16(2), 167-207.