Intermediated Implementation

Anqi Li* Yiqing Xing†

First Draft: December 2015
This Draft: March 2017

Abstract

Many real-world problems such as 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 can only regulate 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 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, whereas regulating the distribution of limited aspects of sold bundles can. We examine the applicability of these results to the regulation of real-world intermediaries.

Keywords: implementation theory; intermediaries; adverse selection; market structure.

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

†Carey Business School, Johns Hopkins University, xingyq@jhu.edu. We thank 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, Anthony Lee Zhang, two anonymous referees, as well as the audience at U of Michigan, UPenn, Stanford, ASSA, ES Summer Meetings and SED for helpful comments and suggestions. Xiaoyu Liu provided excellent and generous graphic 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 particularly when the principal and intermediaries have different objectives. Specifically, we fix any 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 the above described game as a sub-game perfect equilibrium where each agent consumes his bundle under the target social choice rule.

The policies of our interest, namely *per-unit fee schedules* and *distribution regulations*, are commonly used in practice. Our main result shows that the effectiveness of these policies depends 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

[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..

them play the above described game. We show that when intermediaries have private values and are competitive, 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. By contrast, when intermediaries have interdependent values or monopoly power, per-unit fee schedules cannot generally be used to achieve intermediated implementation, whereas regulating the distribution of limited aspects of sold bundles can. We examine the policy implications of these results and discuss their relevance for the regulation of real-world intermediaries.

To illustrate, 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 by 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

[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."

dealers drives profits to zero and sustains the target quality-price rule in *every* subgame perfect equilibrium of the contracting game.

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 from 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 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 achieves 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 as a result of dealers' profitable deviations. Thus in equilibrium, dealers offer the target quality car to each customer, based on which 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, distribution regulation achieves intermediated implementation in *every* equilibrium of the contracting game if and only if permuting target consumption goods among agents does not create new implementable rules. By exploiting the cyclic monotonicity of implementable consumption rules (Rochet (1987)), we show that the above condition is equivalent to (DU), which requires that permuting target consumption goods yields a distinct total utility of consumption among agents. (DU) can be easily satisfied if the type space and the consumption space are single-dimensional and if the agent's utility satisfies the single-crossing property. When the type space or the consumption space is multi-dimensional, the profiles of agent utilities ruled out by (DU) are negligible compared to the set of all feasible utility profiles when the latter is rich.

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 worker's innate ability 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 next 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]

We finally demonstrate the robustness of our results when the principal lacks imprecise knowledge about the underlying environment, and when the technology for serving agents exhibits heterogeneity across intermediaries or non-constant returns to scale. We discuss the relevance of our theory for the regulation of other real-world intermediaries (e.g., financial companies and members of supply chains), as well as its usefulness for guiding empirical studies in related areas.

[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.

1.1 Related Literature

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

Competitive equilibrium under 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 economy. In order to overcome this challenge, many existing studies have impose 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 general equilibrium economies. Miyazaki (1977), Wilson (1977), Riley (1979), Netzer and Scheuer (2014) and the references therein adopt solution concepts where players can anticipate other players' reactions to their contract offers.[5] Last but not least, Guerrieri et al. (2010) demonstrates how search and match frictions are useful for restoring equilibrium existence.

The current paper differs from these studies in several aspects. First, 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 outcomes of laissez-faire economies. Second, we relax the above described restrictions by allowing intermediaries to specify all aspects of the consumption bundle, removing the limit on the menu size, and studying a standard game where intermediaries offer menus of consumption bundles to agents and let agents self-select. Our mechanism for ensuring equilibrium existence differs from that studied by Attar et al. (2011), where latent contracts are proven useful for deterring cream-skimming deviations in the non-exclusive competition for lemons. Stantcheva (2014) examines the optimal redistribution policy in the presence of decentralized

[5]Rothschild and Stiglitz (1976), Wilson (1977) and Riley (1979) limit players to offering one single contract.

employment contracting. However, she adopts Miyazaki (1977)'s solution concept, and this drives the key difference between our results.

Common agency game Our results are stated under the assumption that intermediaries can only offer menus of consumption bundles. That said, notice that under the policies of our interest, all intermediaries offering the full menu of target consumption goods remains an equilibrium of the contracting game even if we extend the contract space to those studied by the existing literature on common agency games (see, e.g., Epstein and Peters (1999), Martimort and Stole (2002) and the follow-up work).

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. 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 problems with quasi-linear utilities and multi-dimensional hidden characteristics, a consumption rule is implementable if and only if it is cyclically monotone. 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 in every equilibrium of the contracting game if any permutation of the target consumption rule is not implementable. This last condition is equivalent to (DU), which requires that any permutation of the target consumption rule changes the total utility of consumption among agents.

Contracting with externality and moral hazard In our setting, each intermediary's contract offer has direct effect on the payoffs 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 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), as the principal there can dictate the full allocation among agents.

The remainder of this paper is organized as follows: Section 2 presents the model setup and illustrative examples; Section 3 analyzes the main results; Section 4 investigates several extensions; Section 5 concludes. See Appendix A for omitted 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.

2.1 Setup

The economy consists of a principal, $I \in \mathbb{N}$ identical intermediaries, and a unit mass of infinitesimal agents whose hidden characteristics denoted by θ are drawn from a compact Borel set $\Theta \subset \mathbb{R}^k$ according to a probability distribution P_θ. A consumption bundle (x, y) consist of two elements x and y that take values in compact Borel sets $X \subset \mathbb{R}^d$ and $Y \subset \mathbb{R}$, where Y is connected in $\mathbb{R}$. Each type θ agent has a unit demand for consumption bundles and derives a utility $u(x, y, \theta)$ from (x, y), whereas a constant-returns-to-scale technology yields a gross profit $\pi(x, y, \theta)$ from serving (x, y) to type θ agent.[6] Throughout, let u and π be continuous and satisfy the following property:

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 zero reservation payoff to players who abstain from the market, i.e., $u(\mathbf{0}, \theta) = 0$ and $\pi(\mathbf{0}, \theta) = 0$ for all θ. 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 agents' type space and the consumption bundle space. For now, suppose that the principal owns the technology for serving agents, and that her goal is to implement a *target social choice rule* $(\hat{x}, \hat{y}) : \Theta \to X \times Y$ that maximizes her unspecified objective function, subject to (1) agents' incentive compatibility constraints

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

and individual rationality constraints, i.e.,

$$u(x(\theta), y(\theta), \theta) \geq u(x(\theta'), y(\theta'), \theta), \forall \theta, \theta' \tag{IC}$$

and

$$u(x(\theta), y(\theta), \theta) \geq 0, \tag{IR}$$

as well as (2) an unspecified feasibility constraint. For the most part, we remain agnostic about the principal's problem and stress that her goal is generally different than maximizing the profit from serving agents. Likewise, we take the target social choice rule as given and focus on the issue of implementation. Interested readers can jump directly to Section 2.2 for concrete examples.

Let $\hat{x}(\Theta)$, $\hat{y}(\Theta)$ and $(\hat{x}, \hat{y})(\Theta)$ denote the image of Θ under mappings $\hat{x}$, $\hat{y}$ and $(\hat{x}, \hat{y})$, respectively. For each $x \in \hat{x}(\Theta)$, let $\hat{y}(x)$ be the unique $y \in \hat{y}(\Theta)$ that satisfies $(x, y) \in (\hat{x}, \hat{y})(\Theta)$,[7] and 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 bundle $(x, \hat{y}(x))$ to its target agents.

Intermediated implementation In the case of *direct implementation* where the principal dictates the menu of consumption bundles, the principal can offer the menu of target consumption bundles $\{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ and let agents self-select. Now suppose the technology for serving agents is owned by intermediaries, whereas the principal can only regulate the x-dimension of sold bundles, due to the various legal, information and administrative barriers that are detailed in Section 2.2. The current analysis takes these barriers as given and studies how the principal can implement the target social choice rule through intermediaries, especially when her goal is different than maximizing intermediaries' profits.

Formally, let $\mathcal{X}$, $\mathcal{Y}$ and Σ denote the Borel-sigma algebra that is restricted to X, Y and Θ, respectively. Let μ_i be the measure on $(X \times Y \times \Theta, \mathcal{X} \otimes \mathcal{Y} \otimes \Sigma)$ that is induced by intermediary i's sold bundles, and let $\mu_{i,x}$ be the measure on $(X, \mathcal{X})$ that is induced by μ_i. Suppose for each $i = 1, \cdots, I$, the principal observes only $\mu_{i,x}$, based

[7] The uniqueness is easy to establish: if there exist $y, y' \in \hat{y}(\Theta)$ where $y < y'$ and $(x, y), (x, y') \in (\hat{x}, \hat{y})(\Theta)$, then all agents prefer (x, y) to (x, y') and hence $(x, y') \notin (\hat{x}, \hat{y})(\Theta)$, a contradiction.

on which she can charge a fee $\psi(\mu_{i,x})$ and leave intermediary i with a net profit of

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

We say that a fee schedule ψ achieves *intermediated implementation* if the following game has a sub-game perfect equilibrium where each agent consumes target consumption bundle:[8]

1. The principal commits to ψ;

2. Intermediaries simultaneously propose menus of *deterministic* consumption bundles, i.e., $\sigma_i \subset 2^{X\times Y}$ for all $i = 1, \cdots, I$;

3. Each agent selects a bundle from $\sigma = \cup_{i=1}^{I}\sigma_i \cup \{\mathbf{0}\}$;

4. Intermediaries deliver the selected bundles to agents and pay fees to the principal.

The above described game best represents situations where intermediaries possess a rich strategy space (i.e., being able to specify all aspects of the bundle and faces no limitation on the size of the menu) whereas agents have no problem locating their most preferred bundles. 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 maximizers.[9]

Two things are noteworthy. First, any player's equilibrium payoff 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 profit $\hat{\pi}(x)$. By taxing this profit away from intermediaries, the principal achieves the exact same outcome as in the case of direct implementation, where she confiscates the profit from selling the target bundle to each type of agent.[10]

[8] It is noteworthy that most of our results hold in every equilibrium of the game.

[9] Our rich strategy space and standard solution concept tend to make it harder, not easier, to establish equilibrium existence. Section 1.1 discusses how existing studies obtain equilibrium existence by adopting different solution concepts or by adding frictions to the contracting game.

[10] Consider a slight variation of our model where the target social choice rule specifies the consumption bundle $(\hat{x}(\theta), \hat{y}(\theta))$ of each agent θ and the net profit $\hat{\Pi}_i$ of each intermediary i. In any equilibrium that achieves intermediated implementation, the principal can first tax away intermediaries' profits and then make a lump sum transfer $\hat{\Pi}_i$ to each intermediary i.

Policies We focus on two kinds of policies that are commonly used in practice (see Section 2.2 for concrete examples). The first kind of policies, hereafter referred to as *per-unit fee schedules*, charges a fee $t(x)$ for every sold bundle (x, y) where $x \in \hat{x}(\Theta)$, as well as a big penalty denoted by "$+\infty$" for every sold bundle (x, y) where $x \notin \hat{x}(\Theta)$.[11] Under such fee schedules, the total charge to intermediary i is

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

The second kind of policies imposes aggregate-level restrictions on the x-dimension of sold bundles. Much attention will be given to *distribution regulations* ψ_{distr}, which implement a fee schedule ψ if the sold bundles match the target probability measure on x, and charge a big penalty otherwise. Formally, let $\hat{P}_x$ denote the probability measure on $(X, \mathcal{X})$ under the target social choice rule, and write the policies of our interest 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 Our 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 payoffs do not depend directly on the agent's 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 perfectly competitive if the above described game involves multiple intermediaries, and it is monopolistic otherwise:

[11]The notation "$+\infty$" means that the penalty is big enough to deter the bundle from being offered on the equilibrium path. 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 expressed as a tax schedule where $T(x) = \begin{cases} \hat{y}(x) & \text{if } x \in \hat{x}(\Theta); \\ +\infty & \text{if } x \notin \hat{x}(\Theta). \end{cases}$

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 hidden tastes θ for the various car models in $X = \{\text{CLA-class, C-class, E-class, etc.}\}$ are denoted by θ. A consumption bundle (x, y) consists of a car model $x \in X$ and a price $y \in Y \subset \mathbb{R}_+$; it incurs a manufacturing cost $c^m(x)$ and yields a utility $u(x, y, \theta) = v(x, \theta) - y$ to type θ customers. The gross profit from distributing a bundle (x, y) to a type θ customer is given by $\pi(x, y, \theta) = y - c^d(x, \theta)$, where the distribution cost $c^d(x, \theta)$ can depend on θ because, e.g., high-taste customers can be more difficult to serve than low-taste customers.

In case the manufacturer owns both the manufacturing technology and the distribution technology, she earns a net profit $\pi(x, y, \theta) - c^m(x)$ from serving a bundle (x, y) to a type θ customer. The target social choice rule maximizes the manufacturer's expected profit subject to customers' incentive compatibility constraints 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 customers. Consider a simple setting where I identical dealers independently decide the price of each car model. By law, the manufacturer cannot influence dealers' pricing strategies or force disclosure of dealers' charged prices or earned profits (see Footnote 1 for Sherman's Act's ban on vertical price fixing). Nevertheless, the manufacturer observes the quantity of each sold car model, based on which she can charge a fee $\psi(\mu_{i,x})$ and leave 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 values 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 bundles and dealers pay fees to the manufacturer. The policies of our interest are commonly observed in reality. In particular,

- $\psi_{per-unit}$ 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 provide dealers with 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 abilities are denoted by θ. An employment contract (x, y) prescribes an after-tax consumption $x \in X \subset \mathbb{R}$ and a 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 generates a profit $\pi(x, y, \theta) = h(y, \theta) - x$ to his employer and the society.

The government wants to implement the social choice rule that maximizes the weighted sum of workers' utilities $\int \lambda(\theta)(x(\theta) - v(y(\theta), \theta))dP_\theta$ subject to workers' (IC) and (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, Mirrlees (1971) suggests that the government propose the menu of target employment contracts $\{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ and let workers self-select; after firms passively execute signed contracts, a typical worker surrenders an income tax $\hat{\pi}(x)$ to the government, retains a consumption x to himself and leaves the firm with zero profit.[12]

But in market economies, firms have the freedom to specify the required performance for each level of pay. And in large and complex organizations, worker performance is best observed by managers and cannot be easily inferred by the government from records that depend on (random) confounding factors (e.g., the company's stock price). The current analysis examines a simple setting where 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

[12] Alternatively, we can write an employment contract $(x + \hat{\pi}(x), y)$ as a pair of pre-tax income $x + \hat{\pi}(x)$ and 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, it suffices to assume that $x/\hat{\pi}(x)$ differs across x.

i, based on which she can charge a tax $\psi(\mu_{i,x})$ and leave firm 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 firms have private values or interdependent values depends on the contracting technology: if y represents effective labor 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)).[13] 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 their most preferred contracts; finally, contracts are executed and the government collects taxes from firms. The policies of our interest are commonly observed in reality. In particular,

- $\psi_{per-unit}$ 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 all x, as 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.[14]

2.2.3 Market-Based Health Care Regulation

A government faces a continuum of patients whose hidden risk types $\theta = (\theta_1, \cdots, \theta_d)$ are drawn from $\Theta = \Delta^d$. Type θ patient's endowment is a random variable that equals $e_s \in \mathbb{R}$ 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 patients and an expected profit $\pi(x, y, \theta) = y - \sum_{s=1}^d \theta_s \cdot (x_s - e_s)$ to the provider.

[13]Section 4.2 discusses the policy implication of this subtle observation.

[14]Hidden disabilities protected by The Americans with Disabilities Act (ADA) include mental illness, traumatic brain injury, HIV/AIDS, diabetes, attention deficit-disorder, etc.. Recent federal initiatives that aim to increase the hiring and retention of employees with disabilities include the Executive Order 13548 signed by President Obama in 2010, which requires that federal agencies reach hiring and retention targets for employees with disabilities.

Under single-payer health care systems, the government proposes to patients 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 patients' (IC) and (IR) and a budget constraint $\int \pi(x(\theta), y(\theta), \theta)dP_\theta \geq 0$. By contrast, under market-based health care systems such as the one adopted by the U.S., many governmental policies regulate some but not all aspects of the insurance policies that are sold by private insurance companies. Due to space limitations, we focus on consumption-plan regulations (e.g., coverage subsidy) in the main body of this paper and defer the analysis of premium regulations (e.g., premium subsidy) to the online appendix.

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. Consider a simple setting where I homogeneous insurance companies independently set 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 $\int_{(x,y,\theta)} \pi(x, y, \theta) - \psi(\mu_{i,x})$.

The current example features interdependent values as the profit from insurance sales directly on the patient's risk type. Time evolves as follows: first, the government commits to a fee schedule; second, insurance companies simultaneously propose menus of insurance products; finally, patients purchase their most preferred products and money changes hands between insurance companies and the government. The policies of our interest have real-life counterparts. In particular,

- $\psi_{per-unit}$ 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

Our main research question concerns which policies of interest can achieve intermediated implementation under what conditions. Our main results are summarized in Table 1, and the message is twofold. First, when intermediaries have private values

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

and are perfectly competitive, intermediated implementation can be achieved through 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 prescribed by the taxation principle in the case of direct implementation, and it allows intermediaries to break even if every $x \in \hat{x}(\Theta)$ is correctly priced at $\hat{y}(x)$; second, it suggests that per-unit fee schedules constitute the *weakest* kind of policies that can potentially achieve intermediated implementation. Under $\hat{\psi}_{per-unit}$, the total charge to intermediary i is

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

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. By contrast, the following aggregate-level distribution regulation — which charges $\hat{\psi}_{per-unit}$ if the intermediary's sold bundles match the target probability measure on x and inflicts a big penalty otherwise — 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 Table 1. Much attention is given to the case of competitive intermediaries in Section 3.1, after which we analyze the case of monopolistic intermediary in Section 3.2 and investigate extensions in Section 4 and the online appendix.

3.1 Competitive Intermediaries

This section examines which policies of our interest achieves intermediated implementation in various environments with competitive intermediaries, with an emphasis on how the answer to this question depends on whether intermediaries have private values and interdependent values.

3.1.1 Private Value and Per-Unit Fee Schedule

Our first theorem shows that when 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; in addition, all sub-game perfect equilibria of the contracting game are payoff equivalent under mild regularity conditions:

Theorem 1. *Suppose that intermediaries have private values and are perfectly competitive and that Assumption 1 holds. Then under* $\hat{\psi}_{per-unit}$*:*

(a) there exists a sub-game perfect equilibrium where $\sigma_i^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ *for all* $i = 1, \cdots, I$*;*

(b) if, in addition, that $\inf_{y \in Y} \pi(x, y) \leq 0$ *for all* $x \in X$*, then all agents obtain their utilities under the target social choice rule and all intermediaries break even in every sub-game perfect equilibrium of the contracting game.*

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

Proof sketch 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 under the target social choice rule and intermediaries have private values, it follows that all bundles in region Π^+ make a profit under $\hat{\psi}_{per-unit}$ regardless of the agent's type. At first sight, it is not obvious what prevents bundles in region Π^+ from being offered or accepted in equilibrium.

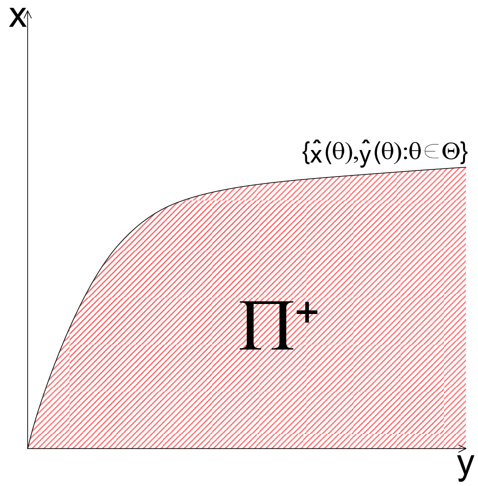

Figure 1: Net profit under $\hat{\psi}_{per-unit}$.

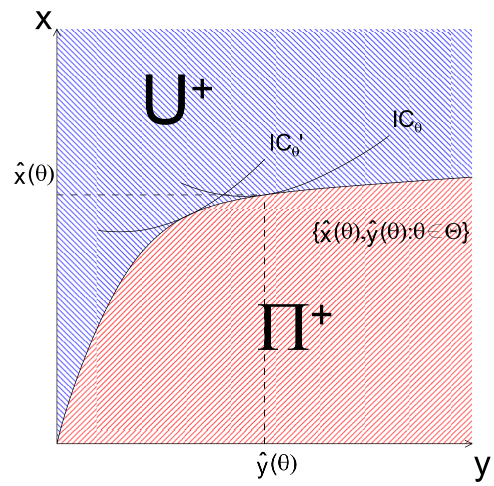

Figure 2: Equilibrium analysis.

Figure 2 adds agents' indifference curves to Figure 1. Since the target social choice rule is incentive compatible, the graph 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 separated, no intermediary can make a profit when the other intermediaries are offering the menu of target consumption bundles, therefore ensuring that all intermediaries offering the full menu of target consumption bundles constitutes an equilibrium outcome of the contracting game. The proof of payoff equivalence is more subtle and is relegated to the appendix.

Alternative interpretation Theorem 1 is closely related to welfare theorems and classical results on price competition. To illustrate, consider an augmented economy where the graph of target consumption bundles represents the intermediary's production-possibility frontier, whereas each agent most prefers his target consumption bundle among all bundles that lie on this frontier. In this economy, the target social choice rule is Pareto-efficient and can thus be sustained as an outcome of the Bertrand competition between intermediaries. A similar observation is made by Fagart (1996) (and is discussed in Section 3.2.1. of Salanié (2005)) in a laissez-faire economy where there is no principal but only profit-maximizing intermediaries.[15]

[15]We thank an anonymous referee for suggesting this reference.

General contract space It is worth noting that while we have so far restricted intermediaries to offering menus of consumption bundles, Part (a) of Theorem 1 nevertheless remains valid even if we extend the contract space to those studied by the existing literature on common agency games (see, e.g., Epstein and Peters (1999), Martimort and Stole (2002) and the follow-up work).

Value-added tax If the principal can regulate intermediaries' profits, too, then intermediated implementation can be achieved by charging intermediaries a one-hundred percent value-added tax. That said, notice that non-governmental entities (e.g., manufacturers) cannot force disclosure of the profit earned by independent entities, and that one hundred percent value-added tax is rarely observed in reality. In case intermediaries already pay a 100α percent value-added tax, intermediated implementation can be achieved by charging a per-unit fee $t(x) = (1-\alpha)\hat{\pi}(x)$ for every sold bundle where $x \in \hat{x}(\Theta)$.

3.1.2 Interdependent Value and Per-unit Fee Schedule

In this section, we illustrate through a simple example why per-unit fee schedules may fail to achieve intermediated implementation when intermediaries have interdependent values and are perfectly competitive:

Example 1. $X \subset \mathbb{R}$, $\Theta = \left\{\underline{\theta}, \overline{\theta}\right\} \subset \mathbb{R}$ and $I \geq 2$. $u(x, y, \theta) = v(x, \theta) - y$ satisfies the single-crossing property (hereafter, SCP) in (x, θ), whereas $\pi\left(x, y, \overline{\theta}\right) > \pi\left(x, y, \underline{\theta}\right)$ for all (x, y). Fix any target social choice rule under which (1) $\hat{y}\left(\underline{\theta}\right) > \inf Y$ and (2) type $\overline{\theta}$ agents' (IC) constraint is binding and type $\underline{\theta}$ agents' (IC) constraint is slack. This last condition, together with SCP, implies that $\hat{x}\left(\underline{\theta}\right) < \hat{x}\left(\overline{\theta}\right)$.

Proposition 1. *No per-unit fee schedule achieves intermediated implementation in the setup of Example 1.*

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 effective labor; finally, per-unit coverage subsidies may fail to induce private insurance companies to provide the desired amount of insurance.

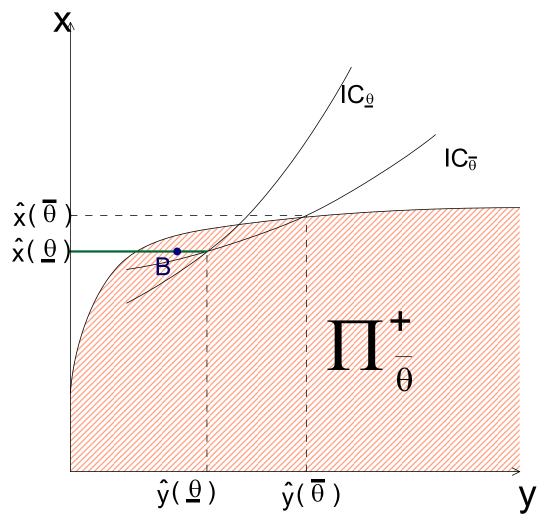

Figure 3: $\hat{\psi}_{per-unit}$ cannot achieve intermediated implementation when intermediaries have interdependent values and are perfectly competitive.

Proof sketch Proposition 1 generalizes the insight of Rothschild and Stiglitz (1976). To gain intuitions, notice first that $\hat{\psi}_{per-unit}$ cannot achieve intermediated implementation in the current example, and the reason is as follows. Now, since intermediaries prefer to serve type $\overline{\theta}$ agents rather than type $\underline{\theta}$ agents, other things being equal, we can no longer separate region $\Pi^{+}_{\overline{\theta}}$ where each bundle makes a profit when it is served to a type $\overline{\theta}$ agent, from region U^{+} where each agent's utility is lower bounded by its counterpart under the target social choice rule. This non-separability is problematic, because in any equilibrium that achieves intermediated implementation, 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, intermediary i's gross profit equals approximately to $\rho\pi\left(\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}), \overline{\theta}\right) + (1-\rho)\pi\left(\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}), \underline{\theta}\right)$, where ρ denotes the population of type $\overline{\theta}$ agents in the economy. In the meantime, the intermediary pays a low fee $\pi\left(\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}), \underline{\theta}\right)$ to the principal and makes a profit overall.[16]

Now take any other per-unit schedule $\psi_{per-unit}$. By definition, in any equilibrium that achieves intermediated implementation under $\psi_{per-unit}$, intermediaries can only break even by incurring a loss from serving a bundle b to its target agents and

[16]Proposition 1 differs from Rothschild and Stiglitz (1976) (hereafter, RS) in two respects. First, we study the implementability of any social choice rule that is incentive compatible, individually rational and feasible among agents, whereas RS examines the outcome of a laissez-faire economy. Second, we allow intermediaries to offer menus of consumption bundles, whereas RS limits insurance companies to offering one single contract.

making a profit from serving a different bundle b' to its target agents, where b and b' are both feasible under the target social choice rule. For concreteness, suppose $b = (\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}))$ and $b' = \left(\hat{x}\left(\overline{\theta}\right), \hat{y}\left(\overline{\theta}\right)\right)$. Take any intermediary i that offers b to type $\underline{\theta}$ agents and notice the following contradiction: on the one hand, if i is the only intermediary who offers b to type $\underline{\theta}$ agents, then all $j \neq i$ offer b' to type $\overline{\theta}$ agents and make a profit, suggesting that i can make a profit, too, by offering $\left(\hat{x}\left(\overline{\theta}\right), \hat{y}\left(\overline{\theta}\right) - \epsilon\right)$ to type $\overline{\theta}$ agents; on the other hand, if another intermediary $j \neq i$ offers b to type $\underline{\theta}$ agents, too, then i can drop b from its menu and save the loss. In the appendix, we complete the argument by showing that no per-unit fee schedule achieves intermediated implementation in the current example.[17]

3.1.3 Distribution Regulation

The negative result of the previous section calls for the use of aggregate-level regulations. In this section, we examine a particular policy in this vein: distribution regulation $\hat{\psi}_{distr}$. We first demonstrate the effectiveness of this policy in a sub-game perfect equilibrium, and then give sufficient and necessary conditions for this policy to achieve intermediated implementation in every sub-game equilibrium.

Throughout this section, let x denote consumption goods and y denote monetary transfers.[18] For the most part, suppose that agents have finite types and that their utilities are quasi-linear in the monetary transfer:[19]

Assumption 2. $|\Theta| = N \in \mathbb{N}$.

Assumption 3. $u(x, y, \theta) = v(x, \theta) - y$.

We first demonstrate the effectiveness of $\hat{\psi}_{distr}$ in a sub-game perfect equilibrium of the contracting game:

Theorem 2. *Suppose that intermediaries are perfectly competitive and that Assumptions 1-3 hold. Then under $\hat{\psi}_{distr}$, there exists a sub-game perfect equilibrium of the contracting game where $\sigma_i^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ for all $i = 1, \cdots, I$ and each agent θ consumes $(\hat{x}(\theta), \hat{y}(\theta))$ on the equilibrium path.*

[17] Miyazaki (1977), Wilson (1977), Riley (1979) and Netzer and Scheuer (2014) use different solution concepts where players can anticipate other players' responses to their contractual offers.

[18] See the online appendix for the opposite case of regulating monetary transfers.

[19] An extension to CARA utilities is investigated at the end of this section.

Thus in reality, regulating the variety of sold products helps achieve the manufacturer's goal whether or not the distribution cost depends on the customer's hidden taste; setting employment targets for big employers restores the effectiveness of Mirrleesian tax schedule whether or not firms can contract directly on the effective labor; finally, Corollary 1 shows that regulating the variety of sold coverage plans helps induce private insurance companies to provide the desired amount of insurance.

Proof sketch Under $\hat{\psi}_{distr}$, intermediaries can only make a profit by permuting target consumption goods across agents. Thus if these deviations are unprofitable for intermediaries, then in equilibrium, all agents purchase their target consumption goods and the competition between intermediaries drives prices to target levels.

Thus the problem boils down to showing that intermediaries cannot profit from permuting target consumption goods across agents. To this end let us first introduce a few definitions. 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 for agents.

To gain intuitions, consider the following example:

Example 2. The economy consists of two types of agents θ_1 and θ_2 who obtain different consumption goods $\hat{x}_1$ and $\hat{x}_2$ under the target social choice rule. By definition, the target consumption rule is implementable and thus satisfies Rochet (1987)'s cyclic monotonicity (CMON), i.e.,

$$v(\hat{x}_1, \theta_1) + v(\hat{x}_2, \theta_2) \geq v(\hat{x}_1, \theta_2) + v(\hat{x}_2, \theta_1).$$

Meanwhile, if permuting target consumption goods across agents yields a new incentive compatible social choice rule that prescribes $(\hat{x}_2, y_2)$ to θ_1 and $(\hat{x}_1, y_1)$ to θ_2, then applying (CMON) to this new rule yields

$$v(\hat{x}_2, \theta_1) + v(\hat{x}_1, \theta_2) \geq v(\hat{x}_1, \theta_2) + v(\hat{x}_2, \theta_1).$$

Combining these conditions shows that any permutation of the target consumption rule is not implementable and thus cannot be part of any profitable deviation for intermediaries if

$$v(\hat{x}_1, \theta_1) + v(\hat{x}_2, \theta_2) \neq v(\hat{x}_1, \theta_2) + v(\hat{x}_2, \theta_1).$$

Meanwhile, if this last condition holds with equality, then agents are indifferent between the bundles prescribed the new social choice rule. By assumption, agents break the tie in favor of their target consumption goods, which in turn prevents the new social choice rule from arising in the first place.

The proof of the general case utilizes the following concepts:

Definition 3. *A bijection $\pi : \Theta \to \Theta$ constitutes a* cyclic permutation of Θ *if $\theta \to \pi(\theta) \to \pi \circ \pi(\theta) \to \cdots$ forms a $|\Theta|$-cycle.*

Definition 4. Permuting a consumption rule $x : \Theta \to X$ yields a distinct total utility of consumption 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*

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

The proof proceeds in two steps. First, we show in the next lemma that if the target 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 a consumption rule $x : \Theta \to X$ 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 of Θ', $x \circ \pi : \Theta' \to X$ is not implementable among agents whose types belong to Θ'.*

Lemma 1 says that intermediaries cannot profit from permuting target consumption goods among agents if the target consumption rule satisfies (DU). Now suppose that the target consumption rule violates (DU). In the appendix, we show that in this case, any unilateral deviation from the equilibrium path is profitable for the intermediary only if it creates cycles where agents swap target consumption goods among each other while being indifferent. Our tie-breaking rule eliminates these cycles and renders the above described deviation unprofitable for intermediaries.

Effectiveness of distribution regulation in every equilibrium The construction so far hinges on agents' off-equilibrium path tie-breaking rule, suggesting that additional assumptions are needed in order to ensure the functionality of distribution regulation in every equilibrium of the contracting game. The next proposition shows that (DU) fulfils this purpose:[20]

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

Proposition 2. *Suppose that intermediaries are perfectly competitive and that Assumptions 1-3 hold. Then under $\hat{\psi}_{distr}$, each agent θ consumes $(\hat{x}(\theta), \hat{y}(\theta))$ in every sub-game perfect equilibrium if and only if Assumption 4 holds.*

The "if" part of Proposition 2 is immediate. To establish the "only if" part, we argue that if the target social choice rule violates (DU), then there exist "bad equilibria" where all intermediaries offer the same menu and compete profits away whereas agents swap target consumption goods among each other on the equilibrium path. Assuming that agents adopt the same tie-breaking rule on and off the equilibrium path, we find that under any unilateral deviation satisfies the distribution requirement, agents must purchase the same consumption goods as they do on the equilibrium path. Thus the deviation is profitable, because profits have already been competed away on the equilibrium path.

In some situations, (DU) is necessary for the functionality of distribution regulation in every sub-game perfect equilibrium of the contracting game even if we do not make assumptions about agents' tie-breaking rules. The next example shows that without (DU), bad equilibria arise not because of agents' off-equilibrium tie-breaking rules, but because of the specificities of intermediaries' profit functions:

Example 2 (Continued)**.** In Example 2, let each agent type constitutes fifty percent of the agent population and let $\pi(x, y, \theta) = y - c(x, \theta)$. Write $u_{ij} = u(\hat{x}_i, \theta_j)$ and $c_{ij} = c(\hat{x}_i, \theta_j)$ for all $i, j = 1, 2$, and normalize $\pi(\hat{x}_i, \hat{y}_i, \theta_i)$ to zero for all $i = 1, 2$.

Suppose that the target consumption rule violates (DU). Consider an outcome that assigns a bundle $(\hat{x}_2, y_2)$ to type θ_1 agents and a bundle $(\hat{x}_1, y_1)$ to type θ_2 agents. In order to sustain this outcome on the equilibrium path, we must have

$$\begin{cases} y_1 - y_2 = u_{12} - u_{22} = u_{11} - u_{21}, \\ y_1 - c_{12} + y_2 - c_{21} = 0, \end{cases}$$

where the first equation says that the outcome is incentive compatible for agents, and

[20]We thank an associated editor who urged us to formalize the necessity of (DU).

the second equation says that intermediaries are breaking even. Denote the solution to this system of equations by (y_1^*, y_2^*).

Now suppose an intermediary unilaterally deviates from the above outcome. In order to make a profit, the deviator must offer $(\hat{x}_1, y_1')$ to type θ_1 agents and $(\hat{x}_2, y_2')$ to type θ_2 agents, where

$$\begin{cases} y_1' - y_2' = u_{12} - u_{22} = u_{11} - u_{21}, \\ y_1' - c_{11} + y_2' - c_{22} \geq 0. \end{cases}$$

In the case where $c_{11} + c_{22} \gg c_{21} + c_{12}$, each type θ_i agents strictly prefer $\left(\hat{x}_{-i}, y_{-i}^*\right)$ to all bundles of the form $(\hat{x}_i, y_i')$, where (y_1', y_2') satisfy the second system of inequalities. Thus the above deviation is unprofitable, a contradiction.

The assumption that intermediaries gain significantly from swapping target consumption goods across agents holds easily when the target social choice rule maximizes the principal's objective but not intermediaries' profits. In the example of taxation, this assumption holds if, e.g., assigning versatile workers to specialized tasks rather than general-purpose tasks boosts profits but hurts positive externalities and hence the social welfare. (DU) rules out these profitable deviations, therefore ensuring that our implementation strategy works independently of intermediaries' profit functions or agents' tie-breaking rules.

(DU) is easy to satisfy in both single-dimensional and multi-dimensional environments:

Example 3. If $X, \Theta \subset \mathbb{R}$ and $v(x, \theta)$ is supermodular in x and θ, then any implementable consumption rule is non-decreasing in the agent's type (Milgrom and Shannon (1994)) and therefore satisfies (DU).

Example 4. Let $X \subset \mathbb{R}^d$ and $\Theta \subset \mathbb{R}^k$, and write $\Theta = \{\theta_1, \cdots, \theta_N\}$. Denote a typical consumption rule by $(x_1, \cdots, x_N)$ where x_i represents the consumption good of agent θ_i. Define $v(x_i, \theta_j) = v_{ij}$ and express the profile of agents utilities by an $N \times N$ matrix V whose ij^{th} entry is v_{ij}. Let $\mathbb{I}$ denote the $N \times N$ diagonal matrix and let Π denote a typical $N \times N$ permutation matrix. By definition, the profiles of agent utilities ruled out by (DU) belong to

$$\{V : \exists\, \Pi \neq \mathbb{I} \text{ s.t. } (\mathbb{I} - \Pi) \cdot V = 0\}.$$

Since this set belongs to a subspace of $\mathbb{R}^{N\times N}$, it is negligible compared to the set of agent utility profiles when the latter is rich.

CARA utility The next corollary extends the analysis so far to CARA utilities and therefore demonstrates the applicability of our results to insurance sales:

Corollary 1. *In the setup of Section 2.2.3, let* $u(x,y,\theta)=\sum_{s=1}^{d}\theta_s\cdot v(x_s-y)$ *where* $v(c)=-\frac{1}{\lambda}\exp(-\lambda c)$*. Then Theorem 2 holds under Assumptions 1 and 2, and Proposition 2 holds under Assumptions 1, 2 and 4.*

Other extensions We examine the effectiveness of coarse distribution regulations in Section 4.1, and extend the analysis to heterogeneous intermediaries and non-constant-returns-to-scale technologies in the online appendix.

3.2 Monopolistic Intermediary

This section examines the case of monopolistic intermediary. Results below hold in both cases of private values and interdependent values.

Per-unit fee schedule Per-unit fee schedules cannot generally be used to achieve intermediated implementation when the intermediary possesses monopoly power. The

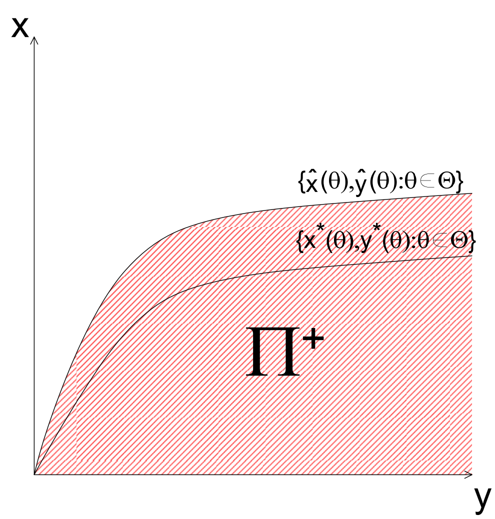

Figure 4: The monopolistic intermediary implements the social choice rule (x^*, y^*) that maximizes its net profit, taking the per-unit fee schedule $\hat{\psi}_{per-unit}$ as given.

reason resembles the logic of monopolistic screening: under most per-unit fee schedules of interest, the monopolistic intermediary can distort the allocations of some agents in order to extract information rents from other agents (see Figure 4 for a graphical illustration).[21]

Distribution regulation As in Section 3.1, let x denote consumption goods and y denote monetary transfers.[22] The assumptions provided below are standard 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. *$u(x, y, \theta) = v(x, \theta) - y$ where $v(x, \theta)$ has SCP in (x, θ) and is continuously differentiable in θ for every x, and $v_\theta(x, \theta)$ is uniformly bounded across (x, θ).*

Define $\hat{P}_y$ and μ_y the same way as we did with $\hat{P}_x$ and μ_x,[23] and denote the $100q^{th}$ percentile of any probability distribution μ over $\mathbb{R}$ by $q(\mu)$. Fix an arbitrary $q \in [0, 1]$ and consider a policy that charges $\hat{\psi}_{per-unit}$ if the intermediary matches the target distribution over consumption goods, and if the $100q^{th}$ percentile of charge prices is upper bounded by its counterpart under the target social choice rule:

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

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

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 housing authorities impose rigid rulings on major real-estate companies regarding the mixture of housing types that can be developed for low-income households. In

[21] The element of monopolistic screening is absent from the literature on double marginalization.

[22] See the online appendix for the opposite case of regulating monetary transfers.

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

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 distributions within the school district.[24]

To better understand Theorem 3, notice that under Assumptions 5 and 6, any implementable consumption rule is non-decreasing in the agent's type (Milgrom and Shannon (1994)). This result, together with the requirement that $\mu_x = \hat{P}_x$, effectively limits the intermediary to offering $x^* = \hat{x}$. Combining this observation with the envelope theorem (Milgrom and Segal (2002)) yields

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

for $\tilde{y} = \hat{y}$ and y^*. Thus the profit-maximizing intermediary will set $y^* = \hat{y}$ in order to meet the requirement $q\left(\mu_y\right) \leq q\left(\hat{P}_y\right)$.

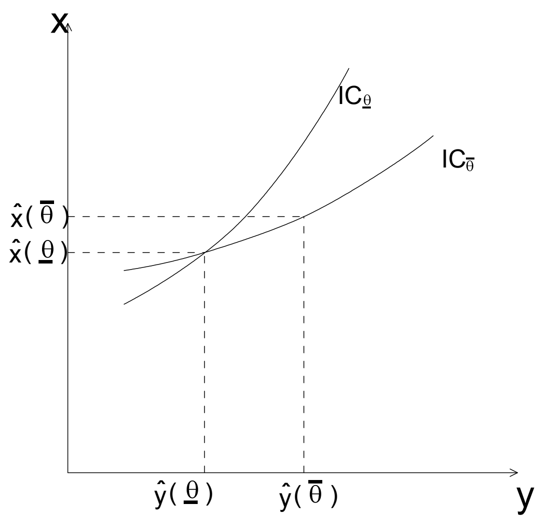

Figure 5: $IC_{\overline{\theta}}$ is binding.

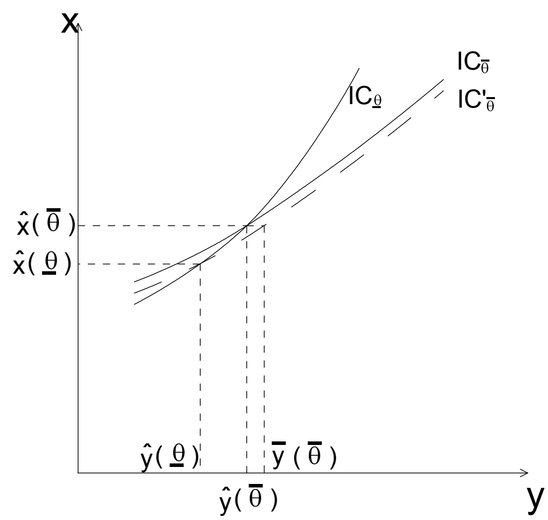

Figure 6: $IC_{\overline{\theta}}$ is slack.

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 incentive constraints bind under the target social choice rule. For example, in the case where $\Theta = \left\{\underline{\theta}, \overline{\theta}\right\}$ and $q = 0$, a close inspection of Figures 5 and 6 reveals that our policy implements the target social

[24]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.

choice rule if and only if $\text{IC}_{\overline{\theta}}$ is binding. When $\text{IC}_{\overline{\theta}}$ is slack, the intermediary offers $\left(\hat{x}\left(\overline{\theta}\right), \overline{y}\left(\overline{\theta}\right)\right)$ to type $\overline{\theta}$ agents instead, and 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, e.g., the distance between target consumption bundles and the elasticity of agents' indifference curves.

4 Discussions

4.1 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 profit function. In order to enforce $\hat{\psi}_{distr}$, she also needs to know the distribution over x under the target social choice rule.

In reality, the principal may not have perfect knowledge about the type distribution P_θ. This limitation has two effects on our results. First, the principal may revise her goal in favor of social choice rules that induce coarse distributions over x. Second, for any given target social choice rule $\hat{x}$, the principal may not be able to enforce the exact distribution requirement $\hat{P}_x$. The analysis below shows how this second problem can be alleviated by coarse distribution regulations.

Formally, suppose that the principal — who knows $\hat{\pi}(x)$ for every $x \in \hat{x}(\Theta)$ — 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 sup-norm):

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

The next corollary shows that the results so far remain valid when ϵ is small:

Corollary 2. *Suppose that intermediaries are perfectly competitive, that Assumptions 1-3 hold and that each agent purchases the bundle that prescribes his target consumption good whenever he is being indifferent. Then under $\hat{\psi}_{distr}$, each type agent θ consumes $(\hat{x}(\theta), \hat{y}(\theta))$ in any sub-game perfect equilibrium of the contracting game when $\epsilon > 0$ is small.*

In Appendix B.2, we demonstrate how coarse distribution regulation enables the principal to virtually implement the target social choice rule through a monopolistic intermediary.

4.2 Contracting Technology

Private value vs. interdependent value We propose two solutions that convert interdependent values into private values. First, notice that intermediaries have private values if parties can contract directly on variables that depend on the agent's type. In the example of income taxation, this means that the government should help improve firms' monitoring technology so that workers can be compensated based on effective labor rather than working hours.

Second, we allow the principal's objective function to depend directly on the agent's type (see Section 2.2 for examples). This observation, albeit a simple one, suggests a second way to eliminate interdependent values: let the principal internalize the part of intermediaries' payoffs that depends directly on the agent's type. In reality, this means that single-payer health care systems do not face the same challenge created by interdependent values as market-based systems do, and that manufacturers should handle maintenance and repair directly if the cost of providing these services varies much with the customer's level of savviness.

Contracting protocol Theorem 1 is independent of parties' contracting protocol: 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 each agent proposes his target consumption bundle in order to maximize his utility while allowing intermediaries to break even. By contrast, this neutrality result breaks down when intermediaries have interdependent values or market power.

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 the various legal, information and administrative barriers. Our research question concerns when and how the principal can implement through intermediaries any target social choice rule that is incentive compatible, individual rational and feasible among agents.

The policies of our interest, namely per-unit fee schedules and distribution regulations, are commonly observed in reality. We show that when intermediaries have private values and are perfectly competitive, the per-unit fee schedule that leaves intermediaries zero profit under the target social choice rule achieves intermediated implementation. By contrast, when intermediaries have interdependent values or monopoly power, per-unit fee schedules generally lose effectiveness, whereas aggregate-level regulations on limited aspects of sold bundles fulfil our purpose. Our proof strategy leverages different characterizations of incentive compatible social choice rules. Applications to sales, taxation and health care regulation are discussed. Extensions to coarse distribution regulations, heterogeneous intermediaries and non-constant-returns-to-scale technologies are finally considered.

We conclude by suggesting potential avenues for future research. First, our analysis is applicable to the regulation of other real-world intermediaries such as financial companies, state-owned enterprises and members of supply chains. In addition, it yields testable predictions about what kinds of regulations we should expect to see, and how the answer to this question depends on the degree of value interdependence and the intermediary's market power. We plan to pursue these issues in future work.

On the theoretical side, it will be interesting to consider other industrial organizations (e.g., horizontal differentiation) and to introduce market frictions into the analysis (e.g., search friction). Our policies prescribe the maximum and the minimum amount of regulations for achieving intermediated implementation in various settings. In the future, we hope to better understand if any in-between policy achieves similar goals, and to characterize the second-best policy when the condition for exact implementation fails.

A Omitted Proofs

Proof of Theorem 1

Proof. Part (a): Fix $\sigma_j^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ for all $j \neq i$. Suppose there exists $(\hat{x}(\theta'), y) \in \sigma_i^*$ such that $u(\hat{x}(\theta'), y, \theta) \geq u(\hat{x}(\theta), \hat{y}(\theta), \theta)$ for some θ, $\pi(\hat{x}(\theta'), y) \geq \psi(\hat{x}(\theta'))$ and one of the these inequalities is strict. Suppose w.l.o.g. that $u(\hat{x}(\theta'), y, \theta) > u(\hat{x}(\theta), \hat{y}(\theta), \theta)$. Since the target social choice rule is incentive compatible, we have $u(\hat{x}(\theta), \hat{y}(\theta), \theta) \geq u(\hat{x}(\theta'), \hat{y}(\theta'), \theta)$ and hence $u(\hat{x}(\theta'), y, \theta) > u(\hat{x}(\theta'), \hat{y}(\theta'), \theta)$.

Under Assumption 1, the last inequality suggests that $y < \hat{y}(\theta')$ and hence that $\psi(\hat{x}(\theta')) \leq \pi(\hat{x}(\theta'), y) < \pi(\hat{x}(\theta'), \hat{y}(\theta')) = \psi(\hat{x}(\theta'))$, a contradiction.

Part (b): Take any equilibrium of the contracting game, where A denotes the set of bundles consumed by agents and $A_i \subset A$ denotes the set of bundles sold by intermediary i. For any set $S \subset A$, let $\mu(S)$ denote the measure of agents who consume bundles in S.

We proceed in two steps. First, suppose, to the contrary, that a positive measure of agents consume bundles that make a profit, i.e., $\sup_{(x,y)\in A} \pi(x,y) \triangleq \overline{\pi} > 0$. Take any $\underline{\pi} \in (0, \overline{\pi})$ and define $B = \{(x,y) \in A : \pi(x,y) \in [\underline{\pi}, \overline{\pi}]\}$. 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 make a 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 consume bundles in $B \cap A_{-i}$, resulting in a gain that is bounded below by $\mu(B \cap A_{-i}) \cdot \underline{\pi}$; second, it undercuts i's own customers and results in 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 any positive measure of sold bundles yields zero profit in every equilibrium of the contracting game.

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

Proof of Proposition 1

Proof. We have shown in Section 3.1.2 that it is impossible to achieve intermediated implementation in any equilibrium where $b = (\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}))$ and $b' = (\hat{x}(\overline{\theta}), \hat{y}(\overline{\theta}))$ (recall that intermediaries incur a loss (resp. make a profit) from selling b (resp. b') to its target agents). Now suppose $b = (\hat{x}(\overline{\theta}), \hat{y}(\overline{\theta}))$ and $b' = (\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}))$. Take any intermediary i who offers b to type $\overline{\theta}$ agents and notice the following contradiction:

on the one hand, if i is the only intermediary who offers b to type $\overline{\theta}$ agents, then all $j \neq i$ offer b' to type $\underline{\theta}$ agents and make a profit, suggesting that i can make a profit, too, by offering $(\hat{x}(\underline{\theta}), \hat{y}(\underline{\theta}) - \epsilon)$ to both types of agents; on the other hand, if some $j \neq i$ offers b to type $\overline{\theta}$ agents, too, i can drop b from its menu and save the loss. $\square$

Proof of Lemma 1

Proof. Take any implementable consumption rule $x : \Theta \to X$ and any $\Theta' \subset \Theta$ where $x(\theta) \neq x(\theta')$ for all $\theta, \theta' \in \Theta'$. Use m to denote the size of Θ', and express any cyclic permutation of Θ' as $\pi(1) \to \pi(2) \to \cdots \to \pi(m) \to \pi(1)$, where π is the induced bijection between $\{1, 2, \cdots, m\}$ and Θ'. For convenience, write $\pi(i) = \theta_i$ for all $i = 1, \cdots, m$.

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 of Θ' yields the result. $\square$

Proof of Theorem 2

Proof. Let $x^i : \Theta \to X$ be the consumption correspondence induced by intermediary i's sold bundles, and let $x^i(\Theta)$ denote the image of Θ under x^i. Apparently, any x^i that meets the distribution requirement satisfies $x^i(\Theta) = \hat{x}(\Theta)$. 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 $\sigma_j^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$ for all $j \neq i$, and 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 generated by (σ_i, σ_j^*). Notice that for each $x \in \hat{x}(\Theta)$, 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, we have $y^i(x) = \hat{y}(x)$, 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 purchase x from $j \neq i$. Thus i cannot meet the distribution requirement and hence is not best-responding.

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

Case (b) This case is ruled out either by (DU) or by agents' off-equilibrium-path tie-breaking rule. 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'', $\cdots$ now. If $\hat{x} : \Theta \to X$ satisfies (DU), then 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 thus cannot arise in the first place. If $\hat{x} : \Theta \to X$ violates (DU), then the inequalities in the proof of Lemma 1 are all binding, suggesting that θ is indifferent between x and x', θ' is indifferent between x' and x'', so on and so forth. Under our tie-breaking rule, the cycle $x \to x' \to \cdots$ cannot arise in the first place.

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

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

2. $x'' \to x'$ implies 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'.

Together, these observations suggest 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'' agents will not purchase $(x', y^i(x'))$, 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 observation, together with $x'' \to x'$, suggests that i is violating the distribution requirement and hence is not best-responding, a contradiction.

Repeating this argument for each x along the chain yields $y^i(x) = \hat{y}(x)$ for all x, from which it follows that $\sigma_i = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$. □

Proof of Proposition 2

Proof. The "if" part: take any equilibrium path of the contracting game, where we use x^i to denote the consumption correspondence induced by intermediary i's sold bundles as before. Clearly, we have $\cup_{i=1}^{I} x^i(\Theta) = \hat{x}(\Theta)$, because otherwise some intermediary is violating the distribution requirement or all intermediaries are abstaining from the market, a contradiction. Now take the graph generated by agents' consumption choices (as in the proof of Theorem 2). Clearly, this graph contains no cycle because of (DU), and it contains no chain because otherwise some intermediary is violating the distribution requirement and prefers to abstain from the market instead. Thus the bundle consumed by each type θ agent takes the form of $(\hat{x}(\theta), y(\theta))$, and the competition between intermediaries drives $y(\theta)$ to $\hat{y}(\theta)$ for all θ.

The "only if" part: suppose the target social choice rule violates (DU). Take any different assignment of target consumption goods than the target consumption rule that satisfies the distribution requirement, and consider the outcome where all intermediaries adopt this assignment rule and compete away profits. Our goal is to sustain this outcome in an equilibrium of the contracting game.

The discussion so far suggests that in the above outcome, an agent either consumes his target consumption good or is indifferent between his target consumption good and someone else's. Assuming that agents adopt this same tie-breaking rule on and off the equilibrium path, it follows that if a unilateral deviation from above outcome satisfies the distribution requirement, then it necessarily assigns the same consumption goods to agents as the above outcome does. But then the deviator cannot make a profit, because profits have already been competed away on the equilibrium path. □

Proof of Corollary 1

Proof. In case agents have CARA utilities, rewrite (IC) as $\tilde{v}(x(\theta),\theta) - \lambda y(\theta) \leq \tilde{v}(x(\theta'),\theta) - \lambda y(\theta')$, where $\tilde{v}(x,\theta) \triangleq \log\left(-\sum_{s=1}^{d} \theta_s \exp(-\lambda x_s)\right)$. Substituting this result into the proofs of Lemma 1, Theorem 2 and Proposition 2 yields the result. □

Proof of Corollary 2

Proof. Take $\epsilon \ll \min_{\theta\in\Theta} P_\theta(\theta)$ and follow the proofs of Theorem 2 and Proposition 2 step by step.

□

B Online Appendix

B.1 Distribution Regulation: Transfers

In the setup of Sections 3.1.3 and 3.2, now let x represent monetary transfers and let y represent 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 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 $\hat{y}(\cdot)$.*

Corollary 3. *Assume 1, 5 and 7. Then:*

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

(ii) suppose that the intermediary is monopolistic and that Assumption 8 holds. Under a slight modification of $\breve{\psi}_{distr}$ where the distribution regulation on y is replaced with $q(\mu_y) = q\left(\hat{P}_y\right)$, we have $\sigma^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \theta \in \Theta\}$.

Proof. Take any incentive compatible social choice rule (x, y). On both sides of the 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 and hence $\hat{x}$ is the only transfer rule that satisfies the distribution requirement.

Part (i): The above discussion suggests that i cannot make a profit and satisfy the distribution requirement simultaneously given σ^*_{-i} and hence 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

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

By Assumption 8, we have $y = \hat{y}$. □

B.2 Coarse Distribution Regulation Against Monopolistic Intermediary

In the setup of Section 3.2, suppose that the p.d.f. of agent types is bounded uniformly below by a constant $f > 0$, and that the target consumption rule $\hat{x}$ is strictly increasing and is c-Lipschitz for some $c > 0$. Denote the c.d.f. of agent types by F. For each $x \in X$, define $\hat{P}_x(x) = F\left(\theta : \hat{x}(\theta) \leq x\right)$, $P^*_x(x) = F\left(\theta : x^*(\theta) \leq x\right)$ and

$P^*_{x,-}(x) = F(\theta : x^*(\theta) < x)$. 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 4. *Suppose that the intermediary is monopolistic and that Assumptions 1, 5 and 6 hold. Then $\|x^* - \hat{x}\| < \frac{c\varepsilon}{f}$ under* (B.1)*.*

Proof. Since $\hat{x}$ is increasing and x^* is non-decreasing, we have

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

and

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

Telescoping yields

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

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

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

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. □

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

In the setup of Section 3.1.3, suppose that when the bundles sold by intermediary i induces a measure μ_i on $(X \times Y \times \Theta, \mathcal{X} \otimes \mathcal{Y} \otimes \Sigma)$, the gross profit of intermediary i is $\pi_i(\mu_i) = \int y d\mu_i - c_i(\mu_{i,x,\theta})$, where $\mu_{i,x,\theta}$ denotes the measure on $(X \times \Theta, \mathcal{X} \otimes \Theta)$ induced by μ_i, and $c_i(\cdot)$ 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 $\hat{\mu}_i(\hat{x}(\theta), \hat{y}(\theta), \theta)$ prescribes the measure of type θ agents who purchase $(\hat{x}(\theta), \hat{y}(\theta))$ from intermediary i 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 agents is served by multiple intermediaries, i.e., for each $\theta \in \Theta$, there exist $i \neq j$ such that $\hat{\mu}_i(\hat{x}(\theta), \hat{y}(\theta), \theta), \hat{\mu}_j(\hat{x}(\theta), \hat{y}(\theta), \theta) > 0$.

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

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

Corollary 5. *Assume 2 and 3. Then under $\{\psi_i\}_{i=1}^I$, there exists a sub-game perfect equilibrium of the contracting game where:*

(a) $\sigma_i^* = \{(\hat{x}(\theta), \hat{y}(\theta)) : \hat{\mu}_i(\hat{x}(\theta), \hat{y}(\theta)) > 0\}$ *for all* $i = 1, \cdots, I$*;*

(b) on the equilibrium path, a measure $\hat{\mu}_i(\hat{x}(\theta), \hat{y}(\theta), \theta)$ *of type* θ *agents purchase* $(\hat{x}(\theta), \hat{y}(\theta))$ *from intermediary* i*.*

Proof. The proof follows that of Theorem 2 step by step. □

References

Attar, A., T. Mariotti and F. Salanié (2011): Non-exclusive Competition in the Market for Lemons, *Econometrica*, 79(6), 1869-1918.

——— (2014): Non-exclusive Competition under Adverse Selection, *Theoretical Economics*, 9(1), 1-40.

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.

Fagart, M. (1996): Concurrence en contrats, anti-sélection et structure d̀information, *Annals of Economics and Statistics*, 43, 1-27.

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.

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.

Salanié, B. (2005): *The Economics of Contracts: A Primer,* MIT Press.

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

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.