

Robustness and Separation in Multidimensional Screening

Gabriel Carroll, Stanford University
gdc@stanford.edu

February 9, 2016

Abstract

A principal wishes to screen an agent along several dimensions of private information simultaneously. The agent has quasilinear preferences that are additively separable across the various components. We consider a robust version of the principal's problem, in which she knows the marginal distribution of each component of the agent's type, but does not know the joint distribution. Any mechanism is evaluated by its worst-case expected profit, over all joint distributions consistent with the known marginals. We show that the optimum for the principal is simply to screen along each component separately. This result does not require any assumptions (such as single-crossing) on the structure of preferences within each component. The proof technique involves a generalization of the concept of virtual values to arbitrary screening problems. Sample applications include monopoly pricing and dynamic taxation.

Keywords: bundling; generalized virtual value; mechanism design; multidimensional screening; robustness; separation.

This paper has greatly benefited from conversations with Florian Scheuer, as well as helpful comments from (in random order) Richard Holden, Dawen Meng, Andy Skrzypacz, Jean Tirole, Alex Wolitzky, Nima Haghpanah, Paul Milgrom, Takuo Sugaya, Drew Fudenberg, Eric Maskin, Ben Brooks, Yiqing Xing, Ilya Segal, Brad Larsen, and Rakesh Vohra, plus conference and seminar audiences at Jiangxi University of Finance and Economics, Becker-Friedman Institute, Simons Institute for Theoretical Computer Science, the University of Western Ontario, and University of Toronto.

1 Introduction

Multidimensional screening stands out among current topics in economic theory as a source of problems that are simple to write down, yet analytically intractable. Whereas the canonical one-dimensional screening model is well-understood, its natural multidimensional analogue appears drastically more complex and unruly.

Consider, for example, the simplest (and most widely studied) version of the problem: A monopolist has two goods, 1 and 2, to sell to a buyer. Marginal costs are zero; the buyer's preferences are quasi-linear and additively separable in the two goods. The monopolist has a prior belief about the joint distribution of the buyer's values for the two goods, (θ^1, θ^2) . She wants to find a mechanism for selling the goods so as to maximize expected revenue. She may simply post a price for each good separately; she may also wish to set a price for the bundle of both goods, which may be greater or less than the sum of the two separate prices. But she can also offer probabilistic combinations of goods, say, a $2/3$ chance of getting good 1 and also a $3/4$ chance of good 2, for yet another price.

The single-good version of the problem is simple: even if probabilistic mechanisms are allowed, the optimum is simply to post a single, take-it-or-leave-it price [33]. Not so in the two-good problem. In the natural case where the values for each good are independent uniform on $[0, 1]$, the optimal mechanism is relatively simple (prices for each good plus a price for the bundle), but known proofs of optimality are involved [29]. In other cases, the optimum may involve probabilistic bundling, and may even require a menu of infinitely many such bundles [43, 17]. Moreover, revenue can be non-monotone — moving the distribution of buyer's types upwards (in the stochastic dominance sense) may decrease optimal revenue [26]. And these phenomena are not artifacts of the possibilities for correlation in the multidimensional problem, since there are examples even when the values for the two goods are independently distributed. Furthermore, if the values are allowed to be correlated, then it can happen that restricting the seller to any fixed number k of bundles can lead to an arbitrarily small fraction of the optimal revenue [25], and finding the optimal mechanism is computationally intractable [16, 18]. With these challenges arising even for the simple monopoly problem, the prospects for more complex multidimensional screening problems are even more daunting.

This paper puts forward an alternative framework for writing models that may escape some of these complexities, and thereby offer new traction. We consider a principal screening an agent with quasi-linear preferences, who has several components of private information. For each component $g = 1, \dots, G$, the agent will be assigned a (possibly

random) allocation x^g , from which he derives a value that depends on his corresponding type θ^g . His total payoff is additively separable across components, $\sum_g u^g(x^g, \theta^g)$. Unlike the traditional model, in which the principal maximizes her expected profit according to a prior distribution over the full type $\theta = (\theta^1, \dots, \theta^G)$, here we assume that the principal only has beliefs about the *marginal* distribution of each component θ^g , but does not know how the various components are correlated with each other. She wishes for a guarantee on her average profit that is robust to this uncertainty. More precisely, she evaluates any possible mechanism according to the worst-case expected profit, over all possible joint distributions that are consistent with the known marginals.

Our main result says that the optimal mechanism separates: for each component g , the principal simply offers the optimal mechanism for that marginal distribution. In the simple monopolist example above, this means that the optimal mechanism is to post a price for each good separately, without any bundling. But the result also allows, for example, that each component g represents a Mussa-Rosen-style price discrimination problem in which different qualities of product may be offered at different prices [31]. In fact, our result is much more general: In the above examples, each component g represents a standard one-dimensional screening problem, satisfying (for example) single-crossing conditions, but we do not actually require this. Ours is a general separability result, in which the structure of preferences within each component can be totally arbitrary. Further applications will be discussed ahead.

Much of the literature on multidimensional mechanism design has emphasized the advantages of bundling (to use the monopolist terminology) — or, more generally, of creating interactions between the different dimensions of private information. In this context, our result expresses the following simple counterpoint: If you don't know enough to be confident you can benefit from bundling, then don't.

Indeed, a simple intuition behind the result is as follows: The solution to a maxmin problem (such as the principal's problem that we set up here) often involves equalizing payoffs across many possible environments. In this case, the separate mechanism does exactly that; since the principal knows the marginal distribution of each θ^g , she can calculate her expected profit, without needing to know anything about how the different components are correlated. However, this intuition is incomplete; hypothetically, there could be some bundling mechanism whose exact profit depends on the correlation structure yet is always better than separate screening.

Our method of proof here works instead by explicitly constructing a joint distribution for which no mechanism performs better than separate screening. The approach first

develops a generalization of the concept of virtual values, applicable to any screening problem; as in traditional one-dimensional problems, a type’s virtual value for an allocation represents the actual value, minus the shadow costs from the incentives of other types to imitate it, and the optimal mechanism assigns each type the allocation for which its virtual value is the highest. Our joint distribution in the multidimensional type space is then chosen so that the (generalized) virtual value of an allocation equals the sum of the virtual values of its components. This condition is expressed as a system of linear equations that the distribution should satisfy. We show that the system indeed has a solution, and we can give this solution at least a mathematical interpretation, if not an economic one; we derive it as the stationary distribution of a particular Markov process over types. A by-product of this proof strategy is our definition of generalized virtual values, which appears to be new to the literature, even though it falls out straightforwardly from the Lagrangian of the principal’s maximization problem.

The main purpose of this of this paper is as a methodological contribution in an area where new tools seem to be called for. Screening problems are now pervasive in many areas of economic theory, and in many applications, agents have private information that cannot be conveniently expressed along a single dimension: In a Mirrlees-style tax model, workers may have multiple dimensions of ability, e.g. ability in different kinds of jobs [35, 36]. In regulation of a monopolist with unknown cost structure [6], it is natural for the costs to be described by more than one parameter, e.g. fixed and marginal costs. In a model of insurance with adverse selection [37], consumers may have different probabilities of facing different kinds of adverse events. In addition, there is a recent growth of interest in dynamic mechanism design, in which information arrives in each period; even if each period’s information is single-dimensional, this field presents some of the same analytical challenges as static multidimensional mechanism design [7]. Rochet and Stole, in their survey [34] (which also describes more applications), put the point more forcefully: “In most cases that we can think of, a multidimensional preference parameterization seems critical to capturing the basic economics of the environment” (p. 150).

This range of applications highlights the need for tractable models for multidimensional screening, to make it easier to gain insight into the economics of these problems. Now, in many applications, the ultimate interest is in interaction across the dimensions of private information, which is exactly what our model here rules out. But this paper offers a very general way of writing down a model where no interactions in the problem formulation lead to no interactions in the solution. This model can then serve as a clean baseline for future explorations where such interactions are systematically added.

Aside from this methodological contribution, our result may also have some value as a positive description of the world. When a buyer walks into a store that sells a thousand different items, she sees separate posted prices for each item (and perhaps special deals for a few combinations of items that are naturally grouped together, or overall bulk discounts). Why has the storekeeper not solved the high-dimensional problem of optimally selling all items simultaneously? One natural answer is that the storekeeper has simply chosen to price items separately as a self-imposed simplicity constraint, to keep her own problem manageable. Our result suggests an alternative take: applying bounded rationality to the structure of the manager’s information instead of the space of mechanisms — perhaps the manager can easily measure the distribution of values for each item separately, but thinking about joint distributions quickly becomes unwieldy, or unreliable due to the curse of dimensionality. Separate pricing then emerges as a natural solution without a priori restricting the space of mechanisms.

This work connects with several branches of literature. It is part of a growing body of work in robust mechanism design, seeking to explain intuitively simple mechanisms as providing guaranteed performance in uncertain environments, and formalizing this intuition by showing how to obtain simple mechanisms as solutions to worst-case optimization problems. This includes earlier work by this author on moral hazard problems in uncertain environments [12, 11], as well as several others, for example [8, 9, 14, 21, 22]. It also relates to a recent spurt of interest in multidimensional screening, particularly in the algorithmic game theory world. Several recent papers in this area have given conditions under which simple mechanisms can be shown to be optimal [23, 45]; most relevantly, the paper [23] introduces a generalized notion of virtual values that is very similar to ours.¹ Others [24, 27, 3, 38] have argued that simple mechanisms for selling multiple goods are approximately optimal under broad conditions. Finally, there is a longstanding, more applied literature on bundling in industrial organization [41, 1, 30, 13], to which we shall return periodically. In contrast to these approximate-optimality and applied strands of literature, the present paper focuses on exact optimality among all possible (deterministic or probabilistic) mechanisms. Still, all of these branches share an interest in understanding or justifying the use of relatively simple mechanisms.

¹The work [10] also defines an object called “virtual types” in the context of a multiple-good, multi-buyer auction, but their virtual types are less close to ours. They play a role in the algorithmic description of the optimal auction, but they do not directly reflect the costs of incentives to other types.

2 Model and result

We shall first lay out the model, and then describe a number of applications.

We will notate the metric on an arbitrary metric space by d ; no confusion should result. For a compact metric space X , we write $\Delta(X)$ with the space of Borel probability distributions on X , with the weak topology.

2.1 Screening problems

We first formally define a *screening problem*, the building block of our model. We assume throughout that there is a single agent with quasi-linear utility.

A *screening environment* (Θ, X, u) is defined by a *type space* Θ and an *allocation space* X , both assumed to be compact metric spaces endowed with the Borel σ -algebra; and a *utility function* $u : X \times \Theta \rightarrow \mathbb{R}$, which is assumed to be (jointly) Lipschitz continuous. We define $Eu : \Delta(X) \times \Theta \rightarrow \mathbb{R}$ as the extension of u by taking expectations over allocations. We will use the variable x to denote either an element of X or of $\Delta(X)$; sometimes we will also use a for elements of X to avoid ambiguity.

In such an environment, the agent may be screened by a *mechanism*. We allow for arbitrary probabilistic mechanisms; thus, a *mechanism* is a pair $M = (x, t)$, with $x : \Theta \rightarrow \Delta(X)$ and $t : \Theta \rightarrow \mathbb{R}$, with t measurable, and satisfying the usual incentive-compatibility (IC) and individual rationality (IR) conditions:²

$$Eu(x(\theta), \theta) - t(\theta) \geq Eu(x(\hat{\theta}), \theta) - t(\hat{\theta}) \quad \text{for all } \theta, \hat{\theta} \in \Theta; \quad (2.1)$$

$$Eu(x(\theta), \theta) - t(\theta) \geq 0 \quad \text{for all } \theta \in \Theta. \quad (2.2)$$

As usual, we are using the revelation principle to justify this formulation of mechanisms as functions of the agent's type; we will generally stick to this formalism, although the verbal descriptions of mechanisms may use other, equivalent formulations.

Note that the formulation here of the IR constraints has assumed each type's outside option is zero. This is essentially without loss of generality, by an additive renormalization of the utility function.

We write \mathcal{M} for the space of all mechanisms.

A *screening problem* consists of a screening environment as above, together with one more ingredient, a given probability distribution over types, $\pi \in \Delta(\Theta)$. Then, the princi-

²Thus we further abuse notation by letting x denote either an allocation or the allocation-rule component of a mechanism; this should not cause confusion in practice.

pal evaluates any mechanism by the resulting expected payment, $E_\pi[t(\theta)]$. We will refer to this objective interchangeably as *revenue* or *profit*.

It is known from the literature (e.g. [5, Theorem 2.2]) that there exists a mechanism maximizing expected payment. Indeed, if Θ is finite then this is easy to see by compactness arguments. The proof for general Θ is much more difficult and will not be presented here. However, for the sake of self-containedness, we may point out that we will not actually need to use the general existence result or its proof in what follows: the statement and proof of our main result, Theorem 2.1, will still be valid using the supremum over mechanisms (without needing to know whether it is attained).

2.2 Joint screening

In our main model, the agent is to be screened according to G pieces of private information, $g = 1, \dots, G$, that enter separately into his utility function. Thus, we assume given a screening problem $(\Theta^g, X^g, u^g, \pi^g)$ for each g . We will refer to these as the *component screening problems*.

These G component screening problems give rise to a *joint screening environment* (Θ, X, u) , where the type space is $\Theta = \times_{g=1}^G \Theta^g$ and the allocation space is $X = \times_{g=1}^G X^g$, with representative elements $\theta = (\theta^1, \dots, \theta^G)$ and $x = (x^1, \dots, x^G)$; and the utility function $u : X \times \Theta \rightarrow \mathbb{R}$ is given by

$$u(x, \theta) = \sum_{g=1}^G u^g(x^g, \theta^g).$$

(Admittedly, there is potential for confusion with the use of these same variables (Θ, X, u) to refer to an arbitrary screening environment in the previous subsection. From now on, they will refer specifically to this product environment except when stated otherwise.)

We will use standard notation such as $\Theta^{-g} = \times_{h \neq g} \Theta^h$; θ^{-g} for a representative element of Θ^{-g} ; and (θ^g, θ^{-g}) for an element of Θ with one component singled out.

The combined utility function u can be extended to probability distributions over X as before. Note however that if ρ is a probability distribution over X , then the expected utility of the agent of type θ is

$$\sum_g E_\rho[u^g(x^g, \theta^g)] = \sum_g E_{\text{marg}_{X^g}(\rho)}[u^g(x^g, \theta^g)],$$

that is, it depends only on the marginal distribution over X^g for each g . Then, we will

think of a random allocation as an element of $\times_g \Delta(X^g)$, and define the expected utility $Eu : \times_g \Delta(X^g) \rightarrow \mathbb{R}$ accordingly. A *mechanism* in this environment then can be defined as a pair of functions (x, t) , with $x : \Theta \rightarrow \times_g \Delta(X^g)$ and $t : \Theta \rightarrow \mathbb{R}$ measurable, satisfying conditions (2.1) and (2.2) as before.

In the joint problem, unlike in Subsection 2.1, we assume that the principal does not know the distribution of the agent's full type θ . Instead, she knows only the marginal distributions of each component, π^1, \dots, π^G , but not how these components are correlated with each other. Our principal possesses non-Bayesian uncertainty about the correlation structure; she evaluates any mechanism (x, t) by its worst-case expected profit over all possible joint distributions. Formally, let Π be the set of all distributions $\pi \in \Delta(\Theta)$ that have marginal π^g on Θ^g for each g . Then, a mechanism (x, t) is evaluated by

$$\inf_{\pi \in \Pi} E_{\pi}[t(\theta)]. \quad (2.3)$$

This describes the joint screening problem. The class of mechanisms we have allowed is fully general: each component x^g of the allocation can depend on the agent's entire type θ . (For example, in selling goods as a bundle, whether the agent receives good 1 depends on his values for the other goods.)

But one thing the principal can always do is to screen each component separately. In particular, let us write R^{*g} for the optimal profit in screening problem $(\Theta^g, X^g, u^g, \pi^g)$, and (x^{*g}, t^{*g}) for the corresponding optimal mechanism. Then define $R^* = \sum_g R^{*g}$. (As mentioned in Subsection 2.1 above, the optimum for each g exists. Even if it did not, we could still take R^{*g} to be the corresponding supremum, and our statement and proof of Theorem 2.1 below would go through essentially unchanged.) The separate-screening mechanism (x^*, t^*) corresponding to these component mechanisms is given by

$$x^*(\theta) = (x^{*1}(\theta^1), \dots, x^{*G}(\theta^G));$$

$$t^*(\theta) = t^{*1}(\theta^1) + \dots + t^{*G}(\theta^G)$$

for each $\theta = (\theta^1, \dots, \theta^G) \in \Theta$. It is immediate that this mechanism satisfies (2.1) and (2.2). And its expected profit is predictable despite the principal's uncertainty: for any possible $\pi \in \Pi$, we always have

$$E_{\pi}[t^*(\theta)] = \sum_g E_{\pi}[t^{*g}(\theta^g)] = \sum_g E_{\pi^g}[t^{*g}(\theta^g)] = \sum_g R^{*g} = R^*.$$

Our main result says that, with the robust objective, separate screening is optimal:

Theorem 2.1. *For any mechanism, the value of the objective (2.3) is at most R^* .*

We show this by constructing a specific distribution π on which the value is at most R^* . More precisely, for most of the proof we work with a discrete approximation (both Θ and X finite), and in that setting we construct a specific π ; we then extend to the continuous case by an approximation argument. Some discussion about possible approaches to constructing π appears in Section 3, followed by the actual proof in Section 4. (A reader eager for the proof can skip ahead to Section 4 without serious loss of continuity.) In addition, a natural follow-up question is how sensitive the result is to the exact joint distribution; we explore this question in Section 5.

2.3 Examples

Here we describe several applications.

Linear monopoly. In the monopoly sales setting described in the introduction, the principal is a seller with G goods available for sale, and the buyer's preferences are additively separable across goods. In this case, θ^g is the buyer's value for getting good g , so Θ^g is an interval in \mathbb{R}^+ , say $\Theta^g = [0, \bar{\theta}^g]$; and $X^g = \{0, 1\}$, with 1 corresponding to receiving the good g and 0 corresponding to not receiving the good. (Remember however that we allow for probabilistic mechanisms, which can give the good with probability between 0 and 1.) The utility function is $u^g(x^g, \theta^g) = \theta^g x^g$.

For each g , π^g is the prior distribution over the buyer's value for good g . In this model, it is well-known (e.g. [33]) that the optimal selling mechanism for a single good is a single posted price p^{*g} : that is $(x^g(\theta^g), t^g(\theta^g)) = (1, p^{*g})$ if $\theta^g \geq p^{*g}$ and $(0, 0)$ otherwise. Consequently, Theorem 2.1 says that, in our model where the seller knows the marginal distribution of values for each good but not the joint distribution of values, the worst-case-optimal mechanism simply consists of posting price p^{*g} for each good g separately.

We shall refer to this example as the *benchmark application*, and shall return to it periodically in discussion.

We have written this application with a single buyer, but note that it works equally well with a continuum population of buyers, in which π^g denotes the cross-sectional distribution of buyers' values for good g . In this case, our restriction to direct mechanisms is not immediately without loss: since the seller does not know the joint distribution π , she could in principle do better by adaptively learning the distribution as in [39], or by

asking each buyer for beliefs about the distribution as in [9, 14]. However, since we actually construct a specific joint distribution that holds the seller down to R^* , this learning would not help her do better in the worst case.

Nonlinear monopoly. More generally, we can consider a multi-good sales model in which each good g can be allocated continuously: $X^g = [0, 1]$. Then we can regard each good g as divisible, and interpret x^g as a quantity; or, alternatively, x^g may be a measure of quality, as in a Mussa-Rosen model [31]. Now θ^g is some parameter describing preferences, and the payoff function $u^g(x^g, \theta^g)$ may be an arbitrary Lipschitz function in our model. Usually in the literature, one imposes structural assumptions on this function — most importantly, an increasing differences condition — that make it possible to solve explicitly for the component- g optimal mechanism using standard techniques. But no such assumptions are needed for our separation result to hold.

In general, the optimal mechanism for component g will now be a menu of various bundles of (possibly probabilistic) quantities and prices (x^g, t^g) , from which each type gets its favorite. Our result then says that, from the point of view of the principal who is uncertain about the joint distribution of the θ^g , the worst-case-optimal joint screening mechanism consists of a separate such menu for each component g .

Although we have implicitly assumed no costs for the principal of producing the good, the model can also easily accommodate costs of production: Suppose that in the component screening problem, the agent’s payoff is given by $u^g(x^g, \theta^g)$, and producing quantity x^g costs $c^g(x^g)$ for the principal; the principal thus wishes to find a mechanism that maximizes $E_{\pi^g}[t^g(\theta^g) - c^g(x^g(\theta^g))]$. This fits into our original model after a renormalization, by defining allocation x^g to be “receiving quantity x^g and paying the production cost” instead of just “receiving quantity x^g .” Explicitly, we apply the original model with payoff function $\tilde{u}^g(x^g, \theta^g) = u^g(x^g, \theta^g) - c^g(x^g)$.

Note also that, while we have insisted on allowing probabilistic mechanisms, there are some sufficient conditions under which the optimal mechanism is actually deterministic. For example, suppose preferences are linear, $u^g(x^g, \theta^g) = \theta^g x^g$, but there is a convex cost function c^g . Then any probabilistic mechanism can be improved on by replacing each type θ^g ’s random allocation by its mean, since this reduces the principal’s cost and has no effect on the IC and IR constraints. Strausz [42] also gives sufficient conditions for deterministic mechanisms to be optimal, although his conditions fall under the “easy case” of our theorem (discussed near the end of Section 3 below), namely when the component problems are standard one-dimensional screening problems in which the monotonicity constraint does not bind.

Optimal taxation. A less obvious application of our model is to a Mirrlees-style taxation problem, in which preferences are quasi-linear and the planner has a Rawlsian objective, to maximize the payoff of the worst-off type. This setup is highly stylized, but it indicates that our framework can help to think about separation in a taxation context.

To illustrate the connection, let us first ignore the joint screening apparatus and return to the general screening language of Subsection 2.1. The taxation problem would be formulated as follows. There is a population of heterogeneous agents, of unit mass, with types indexed by $\theta \in \Theta$ following a known population distribution π . There is a single consumption good. Each agent θ can produce any amount $x \in X = [0, \bar{x}]$ of the good, at a disutility cost $h(x, \theta)$, which we extend linearly to random x (and notate the extension by Eh). We need not make any structural assumptions (e.g. Θ single-dimensional, single-crossing preferences). A mechanism then consists of an allocation rule $x : \Theta \rightarrow \Delta(X)$, and consumption function $c : \Theta \rightarrow \mathbb{R}$, satisfying the incentive constraint

$$c(\theta) - Eh(x(\theta), \theta) \geq c(\hat{\theta}) - Eh(x(\hat{\theta}), \theta) \quad \text{for all } \theta, \hat{\theta}$$

and the resource constraint

$$\int c(\theta) d\pi \leq \int E[x(\theta)] d\pi. \tag{2.4}$$

There is no individual rationality constraint, since everyone can be forced to participate.

The planner's problem is to find a mechanism to maximize the payoff of the worst-off type, $\min_{\theta \in \Theta} (c(\theta) - Eh(x(\theta), \theta))$.

To see how this is equivalent to our formulation, we take the primitive $u(x, \theta)$ to represent the utility an agent would get from producing and consuming x , and the transfer as the net amount redistributed away. Thus, we write $u(x, \theta) = x - h(x, \theta)$. Notice that an allocation rule and consumption rule $(x(\theta), c(\theta))$ then satisfy IC in the taxation problem if and only if the allocation-transfer rule pair $(x(\theta), E[x(\theta)] - c(\theta))$ satisfies IC in the original screening-problem language. Moreover, for any mechanism $(x(\theta), c(\theta))$, adjusting every type's consumption to $c(\theta) + \Delta$ for constant Δ preserves IC, and changes the planner's objective by Δ ; the optimal Δ is the one for which the resource constraint just binds, namely $\int (E[x(\theta)] - c(\theta)) d\pi$. Thus, in the taxation model, the planner's problem is equivalent to maximizing

$$\min_{\theta} (c(\theta) - Eh(x(\theta), \theta)) + \int (E[x(\theta)] - c(\theta)) d\pi$$

over all mechanisms satisfying the IC constraint only. Likewise, in the screening formulation, every type's transfer can be adjusted by a constant Δ , and the optimal Δ is the one that makes IR just bind, namely $\min_{\theta}(c(\theta) - Eh(x(\theta), \theta))$; then, the principal's problem is equivalent to maximizing

$$\int (E[x(\theta)] - c(\theta)) d\pi + \min_{\theta}(c(\theta) - Eh(x(\theta), \theta))$$

over all mechanisms satisfying IC only. So the two problems are equivalent.

Now we move to the joint screening model. Suppose there are multiple income-producing activities $g = 1, \dots, G$; each agent is parameterized by a type for each activity, so the overall type and allocation spaces are $\Theta = \times_g \Theta^g$, $X = \times_g X^g$ with $X^g = [0, \bar{x}^g]$, and payoffs are given by $c - \sum_g h^g(x^g, \theta^g)$. The planner knows the marginal distribution π^g of each θ^g in the population, but not the joint distribution π .

How do we define mechanisms in this model? A mechanism should specify a (probabilistic) level of production in each activity g , and a consumption level, for each type $\theta \in \Theta$, satisfying incentive-compatibility. However it is not clear how the resource constraint should be written when π is unknown. One possible modeling choice would be to allow each agent's allocation and consumption to depend on the entire realized distribution π . Another, much more restrictive, possibility would be to have them depend only on the agent's own type, and require that the resource constraint be satisfied for *every* π , with any surplus resources to be redistributed lump-sum, say.

In any case, one mechanism that will always work is to separate across the activities g : if the optimal tax schedule for activity g is (x^{*g}, c^{*g}) , then each agent θ is assigned to produce $x^{*g}(\theta^g)$ in each activity g , and receive consumption equal to $\sum_g c^{*g}(\theta^g)$. This always satisfies the aggregate resource constraint, for any joint type distribution π . Then, Theorem 2.1 implies that this mechanism is worst-case optimal: no better value of the Rawlsian objective can be guaranteed across all joint distributions (regardless of how we formulate the resource constraint).

If we interpret each g as a time period, then this application of our model connects with a recent literature in dynamic public finance, in which agents' income-producing abilities evolve over time. A prediction of such models is that, in the optimal mechanism, each agent's tax will typically depend on the entire history of his past income. This literature has tacitly acknowledged such history-dependent taxation schemes as unrealistically complicated, and responded by quantitatively comparing with the optima obtained using more restrictive tax instruments, such as ones depending only on age and current income

[19, 40, 20]. However, the theoretical foundations for this approach, or more generally for delineating which kinds of tax systems are or are not “simple,” are yet to be established. Our model gives one avenue for such foundations. In our model, if the planner knows the distribution of ability within each period but not the correlation structure across periods nor the information each agent has about his own future ability, then the optimally robust tax policy will tax and redistribute within each period separately.

3 Unsuccessful proof approaches

As indicated above, we will prove Theorem 2.1 by constructing a particular joint distribution on Θ for which no mechanism can generate profit greater than R^* . In order to better understand the content of the result, we consider some straightforward ways one might try to construct such a joint distribution.

3.1 Independent distributions

One natural first try would be to have the different components θ^g be independently distributed, $\pi = \times_g \pi^g$. A priori, an independent model seems to remove all interactions across components in the setup of the screening problem, so one might expect no interactions in the solution.

However, this approach is a nonstarter even in the benchmark monopoly problem, as is well known from the bundling literature. For example, with a large number of goods with i.i.d. values, the value for the bundle of all goods is approximately pinned down by the law of large numbers; hence, the seller can extract almost all surplus by bundling, which she could not do with separately posted prices [2, 4]. Hart and Nisan [24] show that independence also fails in a minimal example: two goods, where the buyer’s value for each good is either 1 or 2 with probability 1/2 each. The seller can then extract profit 1 for each good, by setting either price 1 or price 2; so her optimal profit from selling the two goods separately is 2. But if the values are independent, she can instead charge a price 3 for the bundle of both goods, which she then sells probability 3/4 and so earns expected profit $9/4 > 2$.

In fact, McAfee, McMillan and Whinston [30] show that with continuous distributions, separate pricing is *never* optimal under independence. This follows from considering the first-order condition for charging a price for the bundle that is just slightly less than the sum of separate prices.

3.2 Maximal positive correlation

Another approach comes from considering the case where the G separate screening problems $(\Theta^g, X^g, u^g, \pi^g)$ are all identical. In this case, one possible joint distribution π is that all components of the agent's type are identical: π is distributed along the diagonal $\{\theta \in \Theta \mid \theta^1 = \dots = \theta^G\}$. For this distribution, the multidimensional joint mechanism design problem is equivalent to the component problem, scaled up by a factor of G ; it is not hard to see that indeed no mechanism can earn expected profit greater than $G \cdot R^* = R^*$, as needed.

How might one generalize this construction when the component problems are not identical? One possibility to try is to have all G components of the type be “as positively correlated as possible,” so as to again reduce the mechanism design problem to a single-dimensional type. For example, in the benchmark monopoly application where $\theta^g \in \Theta^g \subseteq \mathbb{R}$ is the value for good g , let $q^g : [0, 1] \rightarrow \Theta^g$ be the inverse quantile function, defined by

$$q^g(z) = \min\{\theta^g \mid Pr_{\pi^g}(\theta \leq \theta^g) \geq z\},$$

and then define the joint distribution by randomly drawing $z \sim U[0, 1]$ and taking $\theta^g = q^g(z)$ for each g . We refer to this as the *comonotonic* joint distribution.

A problem with this approach is that it is unclear how it would work in general, when each Θ^g is not necessarily single-dimensional. But even in the benchmark application it does not always succeed. For a counterexample, consider two goods. Suppose the possible values for the first good are 1, 2, with probability 1/2 each; and the possible values for the second good are 2, 3, 4, with probabilities 1/3, 1/6, 1/2 respectively. The seller can earn an expected profit of 1 from the first good alone (either by setting price 1 or price 2), and 2 from the second good alone (either by price 2, 3 or 4), so the maximum profit from separate pricing is 3.

In the comonotonic joint distribution π , there are three possible types, (1, 2), (1, 3), (2, 4), occurring with probabilities 1/3, 1/6, 1/2 respectively (as shown in Figure 1(a)). If this is the joint distribution, then we propose the following mechanism: the buyer can either

- receive good 1 at a price of 1;
- receive good 2 at a price of 3;
- receive both goods at a price of 5; or

- receive nothing and pay nothing.

Figure 1(b) shows the regions of buyer space in which each option is chosen; in particular, it is incentive-compatible for the (1, 2)-buyer to buy only good 1, the (1, 3)-buyer to buy only good 2, and the (2, 4)-buyer to buy the bundle, leading to revenue

$$\frac{1}{3} \cdot 1 + \frac{1}{6} \cdot 3 + \frac{1}{2} \cdot 5 = \frac{10}{3} > 3.$$




Figure 1: Counterexample with maximal positive correlation. (a) The candidate joint distribution. (b) A mechanism that outperforms separate sales. (c) A joint distribution for which no mechanism outperforms separate sales.

This shows that the comonotonic distribution cannot be used to prove Theorem 2.1 in general.

Why does the optimal mechanism design problem for this π not decompose into separate problems for each component? One way to think about what goes wrong, expressed in terms of the usual approach to one-dimensional screening problems, is that the monotonicity constraint that arises with multidimensional allocation is weaker than requiring monotonicity good-by-good. This can be seen in the example above: the low and high types both receive good 1, but the middle type does not. This allows the seller to charge a different marginal price for good 1 to the low type than the high type.

Indeed, when the standard solution to the one-dimensional problem has the monotonicity constraint not binding, the maximal-positive-correlation approach does succeed. This is formalized in the following proposition, which for convenience is expressed in terms

of continuous distributions. (The counterexample above is discrete, but this difference is immaterial; it can be made continuous by perturbation.)

Proposition 3.1. *Consider the benchmark monopoly application. Suppose each marginal distribution π^g is represented by a continuous, positive density f^g , and write F^g for the cumulative distribution. Suppose that for each g , there is a type θ^{*g} such that the virtual value*

$$v^g = \theta^g - \frac{1 - F^g(\theta^g)}{f^g(\theta^g)} \quad (3.1)$$

*is negative for $\theta^g < \theta^{*g}$ and positive for $\theta^g > \theta^{*g}$. Let π be the comonotonic joint distribution. For this π , no mechanism yields higher expected profit than R^* .*

This can be shown by the usual method of ignoring the monotonicity constraint, using the infinitesimal downward IC constraints to rewrite profit in terms of virtual surplus, and then maximizing virtual surplus pointwise. The details are in Appendix A. (In fact, Proposition 3.1 extends beyond the benchmark application, to any single-dimensional screening problems with single-crossing in which the monotonicity constraint does not bind; we omit the details but they are straightforward.)

In the light of Proposition 3.1, our Theorem 2.1 can be seen as a result about the relationship between multidimensional ironing and single-dimensional ironing. Indeed, a key object in our proof will be the concept of a *generalized virtual value*, which extends the concept of an ironed virtual value in the benchmark case. After developing this concept, we will show how to use it to directly construct our joint distribution π . In the special case where the component problems involve no ironing, our construction actually recovers the comonotonic distribution (as discussed at the end of Subsection 4.3 below), but in general it is subtler. For the above example, the π we construct is as shown in Figure 1(c).

In this example, there are other joint distributions that also pin down profit to at most R^* ; all of them place positive probability on at least five of the six possible types. However, for the distribution shown, a subset of the constraints — namely, the IC constraints for reporting as the next lower value for one good (and truthful reporting for the other good), plus the IR constraint for the lowest type — are enough to pin down revenue to R^* ; and it is the *unique* distribution with this property.

We turn now to the general proof.

4 The actual proof

We start with a verbal overview of the ideas. The bulk of the proof consists of the case where Θ^g and X^g are finite; afterwards, we extend to the general case by an approximation argument.

For the proof in the finite case, it helps to first develop some basics about the mathematics of screening problems in general. A finite screening problem can be thought of as a linear programming problem, with the components of the mechanism (the probability of each allocation at each type, and the payments) as variables. This allows us to use the language of LP duality, and talk about the dual variables on the IC and IR constraints (also known as Lagrange multipliers).

We first show how, in any finite screening problem, the Lagrangian can be interpreted as a kind of generalized virtual surplus: For each type θ of the agent and each possible allocation x , one can define a generalized virtual value $\bar{u}(x, \theta)$, which represents the effect on revenue from a marginal increase in the probability assigned to allocation x ; it equals the value of type θ for that allocation, minus the shadow cost from binding incentive constraints of other types that would like to imitate type θ . In the familiar case of one-dimensional screening problems, these generalized virtual values are just the usual ironed virtual values; but in fact they can be defined for any finite screening problem. An optimal mechanism then must maximize the virtual surplus for each type separately. Our strategy for proving Theorem 2.1 is to construct the distribution π for the joint problem so that, for each joint type θ , the virtual value for any outcome x is simply the sum of the virtual values $\bar{u}^g(x^g, \theta^g)$ from the component problems. If we can do this, it will immediately follow that the separate-screening mechanism maximizes the virtual surplus and so maximizes revenue.

One complication in this argument is that the generalized virtual values for a screening problem are functions not only of the probability of each type but also of the dual variables. In general, the dual LP may have many optimal solutions, and we need to choose one in order to define the generalized virtual values. (The traditional definition of ironed virtual values implicitly corresponds to one particular choice of dual solution.) Hence, rather than construct a joint distribution π by itself, we will simultaneously construct π and the dual variables for the joint screening problem, using dual variables from the component problems as inputs to the construction.

We can simplify the task of choosing dual variables by using only a subset of the IC constraints — namely, the ones for misreporting a single component of the type; that is,

type $\theta = (\theta^g, \theta^{-g})$ imitating a type $(\hat{\theta}^g, \theta^{-g})$, for some g and $\hat{\theta}^g$. We are thus studying a relaxed problem, and showing that with this subset of constraints, already no mechanism earns more than R^* .

Once we have decided to focus on this set of constraints, it quickly becomes apparent that the generalized virtual values will separate in the desired way if the following relationship holds between the dual variables on the IC's in the joint problem and the corresponding dual variables from the component problems:

$$\lambda[(\theta^g, \theta^{-g}) \rightarrow (\hat{\theta}^g, \theta^{-g})] = \frac{\pi(\theta^g, \theta^{-g})}{\pi^g(\theta^g)} \times \lambda^g[\theta^g \rightarrow \hat{\theta}^g]. \quad (4.1)$$

This formula, in effect, tells us how we should construct the IC dual variables in terms of the joint distribution π (the dual variables on the IR's will be constructed separately). Of course, we cannot simply choose any old π and define the λ 's by (4.1); we need to ensure that they are indeed a solution to the dual LP.

Here is an alternative perspective on what this requirement means. In any screening problem, the dual variables (λ, κ) need to be related to the type probabilities via a certain set of linear equations, which we call *flow equations*. These equations are the constraints dual to the t variables. They also ensure that the principal's profit from any mechanism (the primal objective) is bounded by its generalized virtual surplus.

Our strategy in the joint screening problem, then, is to choose π and the multipliers (λ, κ) so that two conditions hold:

- (i) the flow equations are satisfied;
- (ii) the virtual surplus, when maximized type-by-type, comes out to R^* .

If we can achieve both of these, then together they imply that the profit from any mechanism is at most R^* , which is what we need.

We will choose π , and then determine the multipliers via (4.1) above; this will ensure that virtual values in the joint problem separate into the component virtual values, which will give (ii). It remains to ensure (i). That is, we must show that $\pi \in \Pi$ can be chosen so that the resulting multipliers satisfy the flow equations in the joint screening problem. We will use a connection between the flow equations in the joint problem and the flow equations in the component screening problems to show that this can be done — and to obtain at least something of an interpretation for π in the process.

With this outline to guide us, it's time to begin the details.

4.1 Preliminary tools

We first gather several general-purpose facts about screening problems. For this subsection, we use the variables (Θ, X, u, π) to refer to an arbitrary screening problem as in Subsection 2.1, not the product environment of our main model.

We start with the discussion of generalized virtual values, introduced in the outline above. Suppose that (Θ, X, u, π) is a screening problem in which Θ and X are finite, and every type has positive probability: $\pi(\theta) > 0$ for each $\theta \in \Theta$. (This latter assumption is not without loss of generality: probability-zero types cannot simply be deleted from the type space, because their IR constraints may affect the set of possible mechanisms.)

Let us refer to pure allocations (elements of X) by the variable a . A mechanism then consists of $|\Theta| \cdot (|X| + 1)$ numbers, namely the allocation probabilities $x_a(\theta)$ and payment $t(\theta)$ for each type θ and each $a \in X$, satisfying the IC and IR constraints, as well as the requirement that each $x(\theta)$ form a valid probability distribution:

$$\sum_{a \in X} u(a, \theta)x_a(\theta) - t(\theta) \geq \sum_{a \in X} u(a, \theta)x_a(\hat{\theta}) - t(\hat{\theta}) \quad \text{for all distinct } \theta, \hat{\theta} \in \Theta; \quad (4.2)$$

$$\sum_{a \in X} u(a, \theta)x_a(\theta) - t(\theta) \geq 0 \quad \text{for all } \theta \in \Theta; \quad (4.3)$$

$$x_a(\theta) \geq 0 \quad \text{for all } \theta \in \Theta, a \in X; \quad (4.4)$$

$$\sum_{a \in X} x_a(\theta) = 1 \quad \text{for all } \theta \in \Theta. \quad (4.5)$$

The maximum expected profit over all such mechanisms is given by

$$\max_{(x,t)} \sum_{\theta} \pi(\theta)t(\theta).$$

Maximizing this profit, subject to constraints (4.2)–(4.5), is a linear programming problem.

Let R^* denote the maximal profit. Then, the value of the dual problem is also R^* . Consider an optimal solution to this dual problem. It consists of dual variables for each constraint (these dual variables are indexed by types, which we indicate in square brack-

ets):

$$\begin{aligned}
\lambda[\theta \rightarrow \hat{\theta}] &\geq 0 && \text{for each } \theta \neq \hat{\theta}; \\
\kappa[\theta] &\geq 0 && \text{for each } \theta; \\
\mu_a[\theta] &\geq 0 && \text{for each } \theta, a; \\
\nu[\theta] &&& \text{for each } \theta.
\end{aligned}$$

It will be useful for notational simplicity to also define $\lambda[\theta \rightarrow \theta] = 0$ for each θ .

Writing out explicitly the dual constraints, we have

$$\sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] u(a, \hat{\theta}) - \sum_{\hat{\theta}} \lambda[\theta \rightarrow \hat{\theta}] u(a, \theta) - \kappa[\theta] u(a, \theta) - \mu_a[\theta] - \nu[\theta] = 0 \quad (4.6)$$

for each θ and a , as the constraint corresponding to primal variable $x_a(\theta)$; and

$$\sum_{\hat{\theta}} \lambda[\theta \rightarrow \hat{\theta}] - \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] + \kappa[\theta] = \pi(\theta), \quad (4.7)$$

as the constraint corresponding to $t(\theta)$. We will refer to these latter as *flow equations*, for reasons that will become explicit later. We also have

$$-\sum_{\theta} \nu[\theta] = R^*, \quad (4.8)$$

the optimal value of the objective in the dual.

Now, for each θ , define the *generalized virtual value* for each allocation $a \in X$:

$$\bar{u}(a, \theta) = u(a, \theta) - \sum_{\hat{\theta} \in \Theta} \frac{\lambda[\hat{\theta} \rightarrow \theta]}{\pi(\theta)} \left(u(a, \hat{\theta}) - u(a, \theta) \right), \quad (4.9)$$

and extend this function to randomized allocations, $E\bar{u} : \Delta(X) \times \Theta \rightarrow \mathbb{R}$, by linearity. Also define the maximum generalized virtual value for each type:

$$\bar{u}_{\max}(\theta) = \max_{a \in X} \bar{u}(a, \theta).$$

As we shall momentarily show, these generalized virtual values extend two features of virtual values that are familiar from the traditional setting (e.g. [32]). First, we can use them to eliminate the payments from the principal's maximization problem: for any

mechanism, its profit is at most the (expected) generalized virtual value generated by its allocation rule. (As usual in discrete-type settings, this relation is only a bound, not an equality.) Consequently, we obtain an immediate upper bound on the profit of any mechanism by simply taking the maximum generalized virtual value for each type separately, $\bar{u}_{\max}(\theta)$. Second, any optimal mechanism attains this bound with equality. In particular, for each θ , the optimal mechanism chooses the allocation $a \in X$ with the highest generalized virtual value. (If the optimal mechanism is not deterministic, there must be multiple a 's tied for highest virtual value, and only such a 's can receive any probability weight. Note that this is not reversible: there may exist allocation rules x that put weight only on virtual-value-maximizing outcomes, yet are not part of any optimal mechanism (x, t) .)

To be precise, we will show the following chain of inequalities: for any mechanism (x, t) ,

$$\sum_{\theta} \pi(\theta)t(\theta) \leq \sum_{\theta} \pi(\theta)Eu(x(\theta), \theta) \leq \sum_{\theta} \pi(\theta)\bar{u}_{\max}(\theta) \leq R^*. \quad (4.10)$$

It will then follow immediately that any optimal mechanism (x^*, t^*) satisfies these relations with equality. (In particular, the last inequality is always an equality since it does not depend on the mechanism (x, t) .)

In fact, all of these inequalities are straightforward rearrangements of the duality conditions. For the first inequality in the chain: for any mechanism (x, t) , the IC and IR constraints give

$$\begin{aligned} \sum_{\theta} \pi(\theta)t(\theta) &\leq \sum_{\theta} \left(\pi(\theta)t(\theta) + \sum_{\hat{\theta}} \lambda[\theta \rightarrow \hat{\theta}] \left((Eu(x(\theta), \theta) - t(\theta)) - (Eu(x(\hat{\theta}), \theta) - t(\hat{\theta})) \right) \right. \\ &\quad \left. + \kappa[\theta](Eu(x(\theta), \theta) - t(\theta)) \right). \end{aligned} \quad (4.11)$$

The right-side terms can be reorganized as “ x terms” and “ t terms.” For each θ , the terms containing $x(\theta)$ are

$$\sum_{\hat{\theta}} \lambda[\theta \rightarrow \hat{\theta}]Eu(x(\theta), \theta) - \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta]Eu(x(\theta), \hat{\theta}) + \kappa[\theta]Eu(x(\theta), \theta).$$

Using (4.7) to plug in for $\sum_{\hat{\theta}} \lambda[\theta \rightarrow \hat{\theta}] + \kappa[\theta]$, we see that this equals

$$\left(\sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] + \pi(\theta) \right) Eu(x(\theta), \theta) - \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta]Eu(x(\theta), \hat{\theta}) = \pi(\theta)Eu(x(\theta), \theta).$$

And as for the t terms, we have $t(\theta)$ appearing on the right side of (4.11) with coefficient

$$\pi(\theta) - \sum_{\hat{\theta}} \lambda[\theta \rightarrow \hat{\theta}] + \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] - \kappa[\theta],$$

which is zero by (4.7). This verifies the first inequality in (4.10).

The second inequality is immediate, from the definition of \bar{u}_{\max} .

For the third inequality in (4.10), it suffices to show that *any* function $x : \Theta \rightarrow \Delta(X)$ (without imposing any IC or IR constraints) must satisfy $\sum_{\theta} \pi(\theta) E\bar{u}(x(\theta), \theta) \leq R^*$. For this, we reverse the “ x term” calculations above to obtain

$$\begin{aligned} \sum_{\theta} \pi(\theta) E\bar{u}(x(\theta), \theta) &= \sum_{\theta} \left(\left(\kappa[\theta] + \sum_{\hat{\theta}} \lambda[\theta \rightarrow \hat{\theta}] \right) Eu(x(\theta), \theta) - \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] Eu(x(\theta), \hat{\theta}) \right) \\ &= \sum_{\theta} \sum_a \left(\left(\kappa[\theta] + \sum_{\hat{\theta}} \lambda[\theta \rightarrow \hat{\theta}] \right) u(a, \theta) - \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] u(a, \hat{\theta}) \right) x_a(\theta) \\ &= \sum_{\theta} \sum_a -(\mu_a[\theta] + \nu[\theta]) x_a(\theta) \end{aligned}$$

by (4.6)

$$\leq - \sum_{\theta} \sum_a \nu[\theta] x_a(\theta) = - \sum_{\theta} \nu[\theta] = R^*$$

by (4.8). This completes the argument for (4.10).

At this point, we summarize as a lemma the essential conclusions from the above discussion:

Lemma 4.1. *Consider a screening problem (Θ, X, u, π) in which Θ and X are finite, and $\pi(\theta) > 0$ for each $\theta \in \Theta$. Let R^* be the optimal profit. Then, there exist nonnegative numbers $\lambda[\theta \rightarrow \hat{\theta}]$, for all $\theta, \hat{\theta} \in \Theta$ (with $\lambda[\theta \rightarrow \theta] = 0$), and $\kappa[\theta]$ for all $\theta \in \Theta$, such that*

- *the flow equation (4.7) holds for each θ , and*
- *with generalized virtual values $\bar{u}(a, \theta)$ defined as in (4.9), every possible mechanism (x, t) satisfies the chain of inequalities (4.10), and any optimal mechanism (x^*, t^*) satisfies them with equality.*

As indicated earlier, the generalized virtual values we have defined extend the traditional notion of ironed virtual values from the one-dimensional case. Appendix B works

through this correspondence in detail for the benchmark monopoly problem, showing how the two coincide in this case, so that our definition is a bona fide generalization.

We close this subsection by stating a couple of additional technical results about screening problems that will be useful later. First is a simple continuity result, again for finite screening environments:

Lemma 4.2. *Consider a screening environment (Θ, X, u) , with Θ, X finite. Then, the maximum expected profit is continuous as a function of the distribution π .*

The proof is straightforward, so we leave it to Appendix A.

Second is an approximation lemma due to Madarász and Prat [28] that applies to continuous type spaces. It shows that when each type in a given screening problem is moved by a small amount, the principal's optimal profit does not degrade significantly, even though the optimal mechanism may be discontinuous. The idea is that any given mechanism can be made robust to slight changes in the type distribution by refunding a small fraction of the payment to the agent: doing so pads the incentive constraints in such a way that, if the agent is induced to change his chosen allocation, he does so in favor of more expensive allocations; and this effect outweighs any small change in the agent's location in type space.

Formally, say that two distributions $\pi, \pi' \in \Delta(\Theta)$ are δ -close if Θ can be partitioned into disjoint measurable sets S_1, \dots, S_r such that $d(\theta, \theta') < \delta$ for any θ, θ' in the same cell S_k , and $\pi(S_k) = \pi'(S_k)$ for each S_k .

Lemma 4.3. [28] *Take the environment (X, Θ, u) as fixed. For any $\epsilon > 0$, there exists a number $\delta > 0$ with the following property: For any mechanism (x, t) , there exists a mechanism (\tilde{x}, \tilde{t}) such that*

- (a) *for any two types θ, θ' with $d(\theta, \theta') < \delta$, then $\tilde{t}(\theta') > t(\theta) - \epsilon$;*
- (b) *for any two distributions π, π' that are δ -close,*

$$E_{\pi'}[\tilde{t}(\theta)] > E_{\pi}[t(\theta)] - \epsilon.$$

Again, we include the proof in Appendix A for completeness.

4.2 Main proof: The finite case

Now we start the proof of the theorem proper. We return to the notation $(\Theta^g, X^g, u^g, \pi^g)$ for the component screening problems, and use (Θ, X, u) to refer to the joint environment.

Assume for now that the Θ^g and X^g are all finite.

Also assume that $\pi^g(\theta^g) > 0$ for each g and each type θ^g ; this assumption will be dispensed with at the end of this subsection.

Let $\lambda^g[\theta^g \rightarrow \widehat{\theta}^g]$ and $\kappa^g[\theta^g]$ be dual variables for the component- g problem, given by Lemma 4.1. Thus, they satisfy the flow equations

$$\sum_{\widehat{\theta}^g} \lambda^g[\theta^g \rightarrow \widehat{\theta}^g] - \sum_{\widehat{\theta}^g} \lambda^g[\widehat{\theta}^g \rightarrow \theta^g] + \kappa^g[\theta^g] = \pi^g(\theta^g) \quad (4.12)$$

for each $\theta^g \in \Theta^g$, and with the virtual values defined by

$$\bar{u}^g(x^g, \theta^g) = u^g(x^g, \theta^g) - \sum_{\widehat{\theta}^g} \frac{\lambda^g[\widehat{\theta}^g \rightarrow \theta^g]}{\pi^g(\theta^g)} \left(u^g(x^g, \widehat{\theta}^g) - u^g(x^g, \theta^g) \right),$$

we have from the equality case of (4.10) that

$$\sum_{\theta^g} \pi^g(\theta^g) \bar{u}_{\max}^g(\theta^g) = R^{*g}, \quad \text{where } \bar{u}_{\max}^g(\theta^g) = \max_{x^g} \bar{u}^g(x^g, \theta^g).$$

Also note for future reference that, for each fixed g , if we sum up (4.12) over all θ^g , the $\lambda^g[\dots]$ terms cancel, and we are left with

$$\sum_{\theta^g} \kappa^g[\theta^g] = \sum_{\theta^g} \pi^g(\theta^g) = 1. \quad (4.13)$$

We now begin to construct our joint distribution π and dual variables $\lambda[\theta \rightarrow \widehat{\theta}], \kappa[\theta]$ for the joint problem. We first define the κ 's, by putting

$$\kappa[\theta] = \prod_g \kappa^g[\theta^g].$$

If we multiply (4.13) across all g , we see that

$$\sum_{\theta \in \Theta} \kappa[\theta] = 1. \quad (4.14)$$

And if we fix any g and θ^g , and multiply (4.13) only for all other components $h \neq g$, we likewise get

$$\sum_{\theta^{-g} \in \Theta^{-g}} \kappa[\theta^g, \theta^{-g}] = \kappa^g[\theta^g] \quad (4.15)$$

which will also be useful later.

Next we will define π as the solution to a system of linear equations analogous to (4.12). This system is given by (4.18) below. To briefly motivate it, consider the formula for the generalized virtual value in the joint screening problem: for any $\theta = (\theta^1, \dots, \theta^G)$ and $x = (x^1, \dots, x^G)$,

$$\begin{aligned}\bar{u}(x, \theta) &= u(x, \theta) - \sum_{\hat{\theta} \in \Theta} \frac{\lambda[\hat{\theta} \rightarrow \theta]}{\pi(\theta)} \left(u(x, \hat{\theta}) - u(x, \theta) \right) \\ &= \sum_g \left(u^g(x^g, \theta^g) - \sum_{\hat{\theta} \in \Theta} \frac{\lambda[\hat{\theta} \rightarrow \theta]}{\pi(\theta)} \left(u^g(x^g, \hat{\theta}^g) - u^g(x^g, \theta^g) \right) \right)\end{aligned}$$

by additive separability of u . Our plan is to put positive weight only on IC constraints $\hat{\theta} \rightarrow \theta$ for which $\hat{\theta}, \theta$ differ in just one coordinate. Then, for each g , the difference terms $\left(u^g(x^g, \hat{\theta}^g) - u^g(x^g, \theta^g) \right)$ for which $\hat{\theta}$ and θ differ in a coordinate $h \neq g$ cancel, and we are only left with the terms where they differ in coordinate g ; thus our virtual value will become

$$\sum_g \left(u^g(x^g, \theta^g) - \sum_{\hat{\theta}^g \in \Theta^g} \frac{\lambda[(\hat{\theta}^g, \theta^{-g}) \rightarrow (\theta^g, \theta^{-g})]}{\pi(\theta)} \left(u^g(x^g, \hat{\theta}^g) - u^g(x^g, \theta^g) \right) \right).$$

We wish to choose the π 's and λ 's so that the separate-screening mechanism maximizes generalized virtual value. This will be accomplished if the generalized virtual value of each type θ is the sum of the component virtual values; and for this it is sufficient that, given π , the λ 's should be defined by

$$\lambda[(\hat{\theta}^g, \theta^{-g}) \rightarrow (\theta^g, \theta^{-g})] = \frac{\pi(\theta)}{\pi^g(\theta^g)} \times \lambda^g[\hat{\theta}^g \rightarrow \theta^g]. \quad (4.16)$$

We will also need the joint distribution π and the dual variables $\lambda[\theta \rightarrow \hat{\theta}]$, $\kappa[\theta]$ to satisfy the flow equations

$$\sum_g \sum_{\hat{\theta}^g} \lambda[(\theta^g, \theta^{-g}) \rightarrow (\hat{\theta}^g, \theta^{-g})] - \sum_g \sum_{\hat{\theta}^g} \lambda[(\hat{\theta}^g, \theta^{-g}) \rightarrow (\theta^g, \theta^{-g})] + \kappa[\theta] = \pi(\theta) \quad (4.17)$$

for each $\theta \in \Theta$. Substituting for the λ 's from (4.16), this is equivalent to

$$\sum_g \sum_{\hat{\theta}^g} \frac{\pi(\hat{\theta}^g, \theta^{-g})}{\pi^g(\hat{\theta}^g)} \times \lambda^g[\theta^g \rightarrow \hat{\theta}^g] - \sum_g \sum_{\hat{\theta}^g} \frac{\pi(\theta)}{\pi^g(\theta^g)} \times \lambda^g[\hat{\theta}^g \rightarrow \theta^g] + \kappa[\theta] = \pi(\theta) \quad (4.18)$$

for each $\theta \in \Theta$.

System (4.18) consists of $|\Theta|$ linear equations in $|\Theta|$ unknowns (the probabilities $\pi(\theta)$), whose coefficients come from the component- g problems. The crucial step in our proof is showing that this system indeed has a solution π , which moreover is a probability distribution with the correct marginals.

Lemma 4.4. *Under the assumptions of this subsection (Θ^g finite and $\pi^g(\theta^g) > 0$), there exists a distribution $\pi \in \Pi$ satisfying the flow equations (4.18).*

Proof. First, for each g , we consider a continuous-time Markov process whose state space is Θ^g . (This should not be thought of as an agent's type changing over time; it is simply an abstract process that happens to have this state space.) The process works as follows. There are two types of state changes, both triggered by Poisson arrivals:

- To each type $\hat{\theta}^g \in \Theta^g$ is associated a Poisson clock; when the clock ticks, the state changes to $\hat{\theta}^g$. The intensity of this clock depends on the current state: it is $\lambda^g[\hat{\theta}^g \rightarrow \theta^g]/\pi^g(\theta^g)$ when the current state is θ^g . (The Poisson clocks corresponding to different $\hat{\theta}^g$ run independently conditional on the current state θ^g .)
- In addition, there is an independent “reset” clock at constant Poisson rate 1. When a reset arrives, the type changes to each θ^g with probability $\kappa^g[\theta^g]$. ((4.13) ensures this makes sense — the probabilities add up to 1.)

Let $\rho^g(\theta^g)$ denote the stationary distribution of this Markov process. Because there is a single recurrent set, the stationary distribution is uniquely determined as the solution to the system of linear equations

$$\sum_{\hat{\theta}^g} \rho^g(\hat{\theta}^g) \frac{\lambda^g[\theta^g \rightarrow \hat{\theta}^g]}{\pi^g(\hat{\theta}^g)} + \kappa^g[\theta^g] = \rho^g(\theta^g) \left(\sum_{\hat{\theta}^g} \frac{\lambda^g[\hat{\theta}^g \rightarrow \theta^g]}{\pi^g(\theta^g)} + 1 \right) \quad (4.19)$$

for all $\theta^g \in \Theta^g$. The left-hand side represents the steady-state frequency of changes into state θ^g , and the right-hand side the frequency of changes out of state θ^g . However, from (4.12), we see that distribution π^g is a solution to this system, so it must actually

be the stationary distribution, $\rho^g(\theta^g) = \pi^g(\theta^g)$. (This explains our terminology “flow equations.”)

Now consider a new Markov process whose state space is the joint type space Θ , and which is governed by Poisson state changes as follows:

- For each component g , and each $\hat{\theta}^g \in \Theta^g$, there is a Poisson clock. When it ticks, component g of the state changes to $\hat{\theta}^g$, i.e. if the current state is θ then it changes to $(\hat{\theta}^g, \theta^{-g})$. This clock has an intensity of $\lambda^g[\hat{\theta}^g \rightarrow \theta^g]/\pi^g(\theta^g)$ when the current state is θ .
- In addition, there is a “reset” clock with rate 1; when a reset arrives, the entire state is redrawn, with probabilities $\kappa[\theta]$ for each θ . (This makes sense, by (4.14).)

These clocks are independent across g and $\hat{\theta}^g$, conditional on the current state.

Let π denote the stationary distribution of this Markov process. This stationary distribution must satisfy the equations

$$\sum_g \sum_{\hat{\theta}^g} \pi(\hat{\theta}^g, \theta^{-g}) \frac{\lambda^g[\theta^g \rightarrow \hat{\theta}^g]}{\pi^g(\hat{\theta}^g)} + \kappa[\theta] = \pi(\theta) \left(\sum_g \sum_{\hat{\theta}^g} \frac{\lambda^g[\hat{\theta}^g \rightarrow \theta^g]}{\pi^g(\theta^g)} + 1 \right)$$

for all θ — which is exactly (4.18) above.

Finally, we can see that for each component g of the state, the arrival rate of each type of event that affects component g (either component changes or resets) depends only on the current component g and is otherwise independent of other components. Hence, the evolution of component g of the state is described exactly by our earlier Markov process with state space Θ^g . (This conclusion makes use of (4.15), which ensures that when a reset event occurs, the new θ^g is indeed drawn from distribution κ^g .) Consequently, the stationary distribution π has g -marginal equal to π^g . Thus, $\pi \in \Pi$. \square

Now, armed with this joint distribution π , define the IC multipliers $\lambda[(\hat{\theta}^g, \theta^{-g}) \rightarrow (\theta^g, \theta^{-g})]$ by (4.16) — and write $\lambda[\hat{\theta} \rightarrow \theta] = 0$ whenever $\hat{\theta}, \theta$ differ in more than one component. This ensures (4.17) is satisfied, since it follows from (4.18).

We can use these multipliers to show that no mechanism earns total profit more than R^* against distribution π . The calculation is quite similar to that used for (4.10).

Consider any mechanism (x, t) in the joint screening problem. Because it must satisfy

the IC and IR constraints, we have

$$\begin{aligned}
\sum_{\theta \in \Theta} \pi(\theta) t(\theta) &\leq \sum_{\theta \in \Theta} \left(\pi(\theta) t(\theta) + \sum_{\hat{\theta} \in \Theta} \lambda[\theta \rightarrow \hat{\theta}] \left((Eu(x(\theta), \theta) - t(\theta)) - (Eu(x(\hat{\theta}), \theta) - t(\hat{\theta})) \right) \right. \\
&\quad \left. + \kappa[\theta] (Eu(x(\theta), \theta) - t(\theta)) \right) \\
&= \sum_{\theta} \left(\sum_{\hat{\theta}} \lambda[\theta \rightarrow \hat{\theta}] Eu(x(\theta), \theta) - \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] Eu(x(\theta), \hat{\theta}) + \kappa[\theta] Eu(x(\theta), \theta) \right) \\
&\quad + \sum_{\theta} \left(\pi(\theta) - \sum_{\hat{\theta}} \lambda[\theta \rightarrow \hat{\theta}] + \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] - \kappa[\theta] \right) t(\theta) \tag{4.20}
\end{aligned}$$

just as in (4.10).

Consider the first parenthesized expression on the right, for any given θ (the “ $x(\theta)$ term”). By additive separability of the utility function u , this becomes

$$\sum_g \left(\sum_{\hat{\theta}} \lambda[\theta \rightarrow \hat{\theta}] Eu^g(x^g(\theta), \theta^g) - \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] Eu^g(x^g(\theta), \hat{\theta}^g) + \kappa[\theta] Eu^g(x^g(\theta), \theta^g) \right).$$

(And all λ terms are nonzero only when $\theta, \hat{\theta}$ differ only in one coordinate.) Note that although the payoff separates, we still write $x^g(\theta)$ and not $x^g(\theta^g)$ for the allocation, since x^g can potentially depend on every component of θ . Now using (4.17), this expression rewrites as

$$\begin{aligned}
&\sum_g \left(\left(\pi(\theta) + \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] \right) Eu^g(x^g(\theta), \theta^g) - \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] Eu^g(x^g(\theta), \hat{\theta}^g) \right) \\
&= \sum_g \left(\pi(\theta) Eu^g(x^g(\theta), \theta^g) - \sum_{\hat{\theta}} \lambda[\hat{\theta} \rightarrow \theta] \left(Eu^g(x^g(\theta), \hat{\theta}^g) - Eu^g(x^g(\theta), \theta^g) \right) \right).
\end{aligned}$$

For each $\hat{\theta}$ such that $\hat{\theta}^g = \theta^g$, we have $Eu^g(x^g(\theta), \theta^g) = Eu^g(x^g(\theta), \hat{\theta}^g)$ and the difference term cancels. So the $\lambda[\hat{\theta} \rightarrow \theta]$ terms survive only when $\hat{\theta}$ differs from θ in the

g -coordinate, and we can rewrite our expression as

$$\sum_g \left(\pi(\theta) Eu^g(x^g(\theta), \theta^g) - \sum_{\hat{\theta}^g} \lambda[(\hat{\theta}^g, \theta^{-g}) \rightarrow (\theta^g, \theta^{-g})] \left(Eu^g(x^g(\theta), \hat{\theta}^g) - Eu^g(x^g(\theta), \theta^g) \right) \right).$$

From the definition of the λ 's, (4.16), this equals

$$\begin{aligned} \sum_g \frac{\pi(\theta)}{\pi^g(\theta^g)} & \left(\pi^g(\theta^g) Eu^g(x^g(\theta), \theta^g) - \sum_{\hat{\theta}^g} \lambda^g[\hat{\theta}^g \rightarrow \theta^g] \left(Eu^g(x^g(\theta), \hat{\theta}^g) - Eu^g(x^g(\theta), \theta^g) \right) \right) \\ & = \sum_g \frac{\pi(\theta)}{\pi^g(\theta^g)} \times \pi^g(\theta^g) E\bar{u}^g(x^g(\theta), \theta^g) \\ & = \sum_g \pi(\theta) E\bar{u}^g(x^g(\theta), \theta^g). \end{aligned}$$

Meanwhile, our second parenthesized expression on the right side of (4.20) (the “ $t(\theta)$ term”) equals zero, by (4.17). Therefore, the right side of (4.20) reduces to

$$\sum_{\theta} \sum_g \pi(\theta) E\bar{u}^g(x^g(\theta), \theta^g)$$

and now we can complete the chain:

$$\begin{aligned} \sum_{\theta} \pi(\theta) t(\theta) & \leq \sum_{\theta} \sum_g \pi(\theta) E\bar{u}^g(x^g(\theta), \theta^g) \\ & = \sum_g \left(\sum_{\theta} \pi(\theta) E\bar{u}^g(x^g(\theta), \theta^g) \right) \\ & \leq \sum_g \left(\sum_{\theta} \pi(\theta) \bar{u}_{\max}^g(\theta^g) \right) \\ & = \sum_g \left(\sum_{\theta^g} \left(\sum_{\theta^{-g}} \pi(\theta^g, \theta^{-g}) \right) \bar{u}_{\max}^g(\theta^g) \right) \\ & = \sum_g \left(\sum_{\theta^g} \pi^g(\theta^g) \bar{u}_{\max}^g(\theta^g) \right) \\ & = \sum_g R^{*g} \qquad \qquad \qquad = R^*. \end{aligned}$$

So any joint screening mechanism gives expected profit at most R^* with respect to the

type distribution π , which is what we wanted to show.

This finishes the case where each π^g has full support. To allow for probability-zero types, we use our earlier continuity result, Lemma 4.2.

For each g , suppose π^g is an arbitrary distribution on Θ^g , with R^{*g} the corresponding maximum expected profit. Let π_1^g, π_2^g, \dots be a sequence of full-support distributions on Θ^g that converges to π^g . Then, if we let $R_1^{*g}, R_2^{*g}, \dots$ be the corresponding values of the maximum expected profit, Lemma 4.2 says that $R_n^{*g} \rightarrow R^{*g}$ as $n \rightarrow \infty$.

For each n , the proof we have just completed shows that there exists a joint distribution π_n on Θ , with marginals π_n^g , such that no mechanism earns an expected profit of more than $R_n^* = \sum_g R_n^{*g}$ with respect to π_n . By compactness, we can assume (taking a subsequence if necessary) that π_n converges to some distribution π on Θ . Then π has marginal π^g on each Θ^g . And for any mechanism (x, t) , we have $E_{\pi_n}[t(\theta)] \leq R_n^*$; taking limits as $n \rightarrow \infty$, we have $E_\pi[t(\theta)] \leq R^*$. Thus, no mechanism earns expected profit more than R^* on distribution π .

4.3 Main proof: The general case

Now we can prove Theorem 2.1 in general, with potentially infinite type and allocation spaces.

First, extending to general allocation spaces, but keeping type spaces finite, is completely straightforward. Indeed, suppose each Θ^g is finite, and let R^{*g} be the optimal revenue in component g , and $R^* = \sum_g R^{*g}$. Suppose there exists a mechanism (x', t') that achieves revenue at least $R^* + \epsilon$ with respect to every joint distribution $\pi \in \Pi$, where $\epsilon > 0$.

For each component g , define an auxiliary screening problem $(\tilde{\Theta}^g, \tilde{X}^g, \tilde{u}^g, \tilde{\pi}^g)$ as follows: $\tilde{\Theta}^g = \Theta^g$ and $\tilde{\pi}^g = \pi^g$; $\tilde{X}^g = \{x'^g(\theta) \mid \theta \in \Theta\}$; $\tilde{u}^g(x^g, \theta^g) = Eu^g(x^g, \theta^g)$. That is, we consider only the component- g allocations which were actually assigned to some type in the mechanism (x', t') ; such an allocation may have been a lottery, but we treat it as a pure allocation in the new screening problem. It is evident that any mechanism for this new screening problem translates to a mechanism for the original component- g screening problem, with the same expected profit; hence, the optimal profit for the new component- g screening problem is $\tilde{R}^{*g} \leq R^{*g}$. Because $\tilde{\Theta}^g$ and \tilde{X}^g are finite, we can apply the case of the previous subsection to see that there is a joint distribution $\tilde{\pi}$ for which no joint screening mechanism can give expected profit more than R^* . But (x', t') evidently gives us a joint screening mechanism for the new problem, which produces expected profit $\geq R^* + \epsilon$

for every possible joint distribution — a contradiction.

Finally, we can remove the restriction to finite type spaces. Here is where we will apply Lemma 4.3, the approximation result for continuous types. Consider the general setting of Subsection 2.2. Let R^* be defined as in that section, and suppose, for contradiction, that there exists a mechanism (x', t') that achieves profit at least $R^* + \epsilon$ against every possible joint distribution π , where $\epsilon > 0$.

For each component g , let δ^g be as given by Lemma 4.3 with $\epsilon/(2G+1)$ as the allowable profit loss. By making δ^g smaller if necessary, we may also assume that any two types at distance at most δ^g have their utility for every $x^g \in X^g$ differ by at most $\epsilon/(2G+1)$. Also, consider the joint screening environment, with the metric on Θ given by the sum of componentwise distances, and let δ be as given by Lemma 4.3 with $\epsilon/(2G+1)$ as allowable loss again. Now let $\bar{\delta} = \min\{\delta^1, \dots, \delta^G, \delta/G\}$.

For each g , we form an approximate distribution on Θ^g with finite support, as follows: Let $\tilde{\Theta}^g$ be a finite subset of Θ^g , with the property that every element of Θ^g is within distance $\bar{\delta}$ of some element of $\tilde{\Theta}^g$ (this can be done by compactness). Arbitrarily partition Θ^g into disjoint measurable subsets $S_{\theta^g}^g$, for $\theta^g \in \tilde{\Theta}^g$, such that each element of any $S_{\theta^g}^g$ is within distance $\bar{\delta}$ of θ^g , and θ^g itself is in $S_{\theta^g}^g$. Then define a distribution $\tilde{\pi}^g \in \Delta(\Theta^g)$, supported on $\tilde{\Theta}^g$, by $\tilde{\pi}^g(\theta^g) = \pi^g(S_{\theta^g}^g)$.

Evidently, $\tilde{\pi}^g$ is $\bar{\delta}$ -close to π^g . Therefore, the maximum profit attainable in the screening problem $(\Theta^g, X^g, u^g, \tilde{\pi}^g)$ is at most $R^{*g} + \epsilon/(2G+1)$: otherwise, Lemma 4.3 would be violated in going from $\tilde{\pi}^g$ from π^g .

In turn, we can view $\tilde{\pi}^g$ as a distribution just on $\tilde{\Theta}^g$, and any mechanism for screening problem $(\tilde{\Theta}^g, X^g, u^g, \tilde{\pi}^g)$ can be converted to a mechanism on the whole type space Θ^g , by assigning each type its preferred outcome from the set $\{(x^g(\theta^g), t^g(\theta^g)) \mid \theta^g \in \tilde{\Theta}^g\}$, and subtracting $\epsilon/(2G+1)$ from all payments (in order to make sure that IR is still satisfied for types outside of $\tilde{\Theta}^g$). This conversion causes a profit loss of $\epsilon/(2G+1)$. We conclude that the maximum profit attainable in screening problem $(\tilde{\Theta}^g, X^g, u^g, \tilde{\pi}^g)$ is at most $R^{*g} + 2\epsilon/(2G+1)$.

Since this is true for each g , we can apply the finite-type-space version of our result to conclude the following: there is a distribution $\tilde{\pi}$ on $\tilde{\Theta} = \times_g \tilde{\Theta}^g$, with marginals $\tilde{\pi}^g$, for which any mechanism earns expected profit at most $R^* + 2G\epsilon/(2G+1)$.

Now we will construct a measure on Θ based on our discretization, but with marginals given by the original π^g . For each $\theta = (\theta^1, \dots, \theta^G) \in \tilde{\Theta}$, let $S_\theta \subseteq \Theta$ be the product set $\times_g S_{\theta^g}^g$; notice that as θ varies over $\tilde{\Theta}$, the sets S_θ form a partition of Θ .

For each such θ , we define a measure π_θ on S_θ as follows. If $\tilde{\pi}^g(\theta^g) = 0$ for some g , let

π_θ be the zero measure. Otherwise, consider the conditional probability measure $\pi^g|_{S_{\theta^g}^g}$ for each component g (which is well-defined since π^g assigns positive probability to $S_{\theta^g}^g$). Define π_θ to be the product of these conditional measures, multiplied by the scalar $\tilde{\pi}(\theta)$.

We can extend π_θ to a measure on all of Θ (by taking it to be zero outside S_θ). Now simply define a measure π on Θ as the sum of π_θ , over all $\theta \in \tilde{\Theta}$.

Note that π is indeed a probability measure; this follows from the fact that the measure assigned to any S_θ equals $\tilde{\pi}(\theta)$, and hence the total measure of Θ is

$$\pi(\Theta) = \sum_{\theta \in \tilde{\Theta}} \pi(S_\theta) = \sum_{\theta \in \tilde{\Theta}} \tilde{\pi}(\theta) = 1.$$

In fact, the marginal of π along component g must equal π^g , for each g . This follows from two facts: first, the probability assigned to any cell $S_{\theta^g}^g$ under this marginal is equal to

$$\sum_{\theta^{-g} \in \tilde{\Theta}^{-g}} \pi(S_{(\theta^g, \theta^{-g})}) = \sum_{\theta^{-g} \in \tilde{\Theta}^{-g}} \tilde{\pi}(\theta^g, \theta^{-g}) = \tilde{\pi}^g(\theta^g) = \pi^g(S_{\theta^g}^g);$$

and second, conditional on cell $S_{\theta^g}^g$, the distribution along the g -component follows $\pi^g|_{S_{\theta^g}^g}$ (because this is true for each of the cells $S_{(\theta^g, \theta^{-g})}$).

Consequently, our mechanism (x', t') satisfies $E_\pi[t'(\theta)] \geq R^* + \epsilon$, by assumption.

In addition, the fact that $\pi(S_\theta) = \tilde{\pi}(\theta)$ for every $\theta \in \tilde{\Theta}$ implies that $\tilde{\pi}$ is δ -close to π , since every element of $S_{\theta^g}^g$ is within distance $\bar{\delta} \leq \delta/G$ of θ^g , for each component g . Consequently, Lemma 4.3 assures us the existence of a mechanism whose expected profit with respect to $\tilde{\pi}$ is greater than $(R^* + \epsilon) - \epsilon/(2G + 1) = R^* + 2G\epsilon/(2G + 1)$.

But we already showed it is impossible to earn profit greater than $R^* + 2G\epsilon/(2G + 1)$ against $\tilde{\pi}$. Contradiction. □

At this point, having completed the proof of the theorem, we can briefly remark on the construction of π embedded in Lemma 4.4, and informally relate it back to the attempts in Section 3. Consider the benchmark monopoly application, and suppose the marginal distributions π^g have increasing virtual values. Take a fine discretization of the type space. For each good g separately, the standard solution to the dual problem puts positive weight only on the adjacent downward incentive constraints and the IR constraint of the lowest type. (This dual solution is constructed explicitly in Appendix B.) Thus, the component- g Markov process in the proof of Lemma 4.4 runs as follows: the state θ^g gradually increases from one type to the next higher type according to a Poisson clock,

interrupted by resets at which the state moves back to the lowest type. In the continuous-type limit, the state drifts deterministically upward in the type space, until a reset occurs, when the process returns to the lowest type. Accordingly, the corresponding joint type process runs as follows: in between resets, the state θ drifts deterministically upward in every component; when a reset occurs, all components jump back to their lowest possible value. Consequently the movements of all components are perfectly synchronized, so the stationary distribution must be the comonotonic distribution that we saw in Section 3. (However, for any finite approximation of the type space, the components are not perfectly synchronized, and the resulting joint distribution π has full support.)

5 Sensitivity analysis

A natural response to Theorem 2.1 is to ask how sensitive the result is to the assumption of extreme uncertainty about the joint distribution π . In particular, our proof approach constructs a very specific π — and one that does not necessarily have an obvious economic meaning palatable for applications (unlike, say, the independent distribution). If there were some other mechanism that performed better than separate screening as long as the joint distribution is not this specific π , that would lessen the appeal of the result — both its methodological appeal (since it would be a strike against the reasonableness of the worst-case formulation), and its literal appeal as an explanation of, say, separate posted prices in the real world. Accordingly, this section will briefly attempt to investigate whether separate screening ceases to be optimal once the designer has a little information about the joint distribution. For concreteness, we focus on the benchmark monopoly application throughout this section.

One immediate difficulty is that it is not clear exactly what it means to have “a little information” about π . One might impose some one-dimensional moment restriction, of the form

$$E_\pi[z(\theta)] = 0 \tag{5.1}$$

or

$$E_\pi[z(\theta)] \geq 0 \tag{5.2}$$

for some function $z : \Theta \rightarrow \mathbb{R}$, and let Π' be the set of all $\pi \in \Pi$ satisfying the restriction, and then evaluate a mechanism (x, t) according to $\inf_{\pi \in \Pi'} E_\pi[t(\theta)]$ instead of the original objective (2.3), and ask whether separate sales is still optimal. However, it is clear that the answer will sometimes be negative. For example, we know that some other mechanisms

(x', t') will *sometimes* give higher profit than R^* , so if our moment restriction is of the form $E_\pi[t'(\theta)] \geq R'$ for $R' > R^*$, clearly our result fails. The question then seems to be, what are the interesting restrictions to consider?

We explore two approaches here. First, we consider one particular moment restriction suggested by the literature, namely negative correlation in values between two goods, and explore its consequences numerically. Second, we show that in the finite-type version of the model, in general there is an open set of distributions π for which separate pricing is optimal; this gives some assurance that Theorem 2.1 is not a knife-edge result, without needing to take a stand on specific moment restrictions.

5.1 Negative correlation

One of the main intuitions from the early bundling literature is that bundling is profitable when values for goods are negatively correlated. This argument has been passed down as received wisdom (see e.g. the undergraduate text by Church and Ware [15, pp. 169–170]), although it comes chiefly from examples rather than any general theorem [41, 1]: for example, in the extreme case where the total value for all goods is deterministic, the seller can extract the full surplus by selling the bundle of all goods. It is also in some sense opposite to the positively-correlated case of Subsection 3.2. For both these reasons, negative correlation seems like a natural place to look for restrictions that might overturn the separation result.

Accordingly, let us take the restriction of negative correlation literally, and test it by considering $G = 2$ goods and imposing the moment restriction

$$E_\pi[\theta^1 \theta^2] \leq E_{\pi^1}[\theta^1] \times E_{\pi^2}[\theta^2]$$

on the possible joint distributions $\pi \in \Pi'$. It turns out that with this restriction, it may or may not happen that the worst-case-optimal revenue is still the R^* from selling separately.

To get some sense of whether one case or the other is common, random numerical experiments were run in Matlab, using finite type spaces Θ . Note that for any one-dimensional moment restriction of the form $E_\pi[z(\theta)] \geq 0$, the worst-case-optimization problem

$$\max_{(x,t) \in \mathcal{M}} \left(\min_{\pi \in \Pi'} E_\pi[t(\theta)] \right) \tag{5.3}$$

can be computationally implemented as follows: Since the inner minimization is a linear program (with the probabilities $\pi(\theta)$ as the choice variables), by LP duality, it can also

be expressed as a maximization problem, namely

$$\max_{\alpha^g[\theta^g], \beta} \sum_g \sum_{\theta^g} \pi^g(\theta^g) \alpha^g[\theta^g]$$

over all choices of real numbers $\alpha^g[\theta^g]$ (for each $g = 1, \dots, G$ and $\theta^g \in \Theta^g$) and $\beta \geq 0$ satisfying

$$\sum_g \alpha^g[\theta^g] + \beta z(\theta) \leq t(\theta) \quad \text{for all } \theta = (\theta^1, \dots, \theta^G) \in \Theta. \quad (5.4)$$

Therefore, the problem in (5.3) can be expressed as a single maximization — over all choices of the mechanism (x, t) , and all $\alpha^g[\theta^g]$ and β satisfying (5.4) — that can be solved as a standard LP.

In the simulations, we randomly generated 1000 choices for the marginal distributions (π^1, π^2) , calculated the mechanism that maximizes the worst-case revenue over negatively correlated distributions as just described, and checked whether the worst-case revenue was strictly higher than obtained by selling the two goods separately. The set of possible values for each good was $\Theta^1 = \Theta^2 = \{1, 2, 3, 4, 5\}$, and the marginals π^1, π^2 were generated by drawing a probability uniformly from $[0, 1]$ for each value, then rescaling to make the probabilities sum to 1.

The result was that in 975 of the 1000 trials, the worst-case revenue was still the R^* from selling separately. That is, based on the simulation results, the optimality of separate sales is usually robust to assuming that the values are known to be negatively correlated.

One might protest that this sensitivity test is too weak, because it is inappropriate to take negative correlation so literally; it is no surprise that one misspecified inequality restriction often fails to rule out some worst-case distributions. It is not clear what an appropriate sharper test would be, but one possible test would be to impose negative *affiliation* on π — that is, negative correlation conditional on $\theta \in \tilde{\Theta}^1 \times \tilde{\Theta}^2$, for all nonempty measurable subsets $\tilde{\Theta}^1 \subseteq \Theta^1, \tilde{\Theta}^2 \subseteq \Theta^2$. This limits the possible π to a much smaller set, so one would expect separate sales to be worst-case optimal much less often.

It is not immediately clear how to repeat the above computational exercise with the restriction of negative affiliation, since it is not a linear constraint on π , so the worst-case optimization cannot readily be expressed as a linear program. However, if we move to the continuous-type setting, we can obtain a negative result: separate sales is (essentially) *always* dominated by bundling. The dominance is weak for a given bundle price, but can be made strict by randomizing the bundle price.

Proposition 5.1. *Consider the benchmark monopoly application, with $G = 2$. Suppose each set of values Θ^g is an interval in \mathbb{R}^{++} , and each marginal distribution π^g is represented by a continuous, positive density f^g . Assume that for each g , the optimal separate price p^{*g} is strictly in the interior of Θ^g . Then there exist $\underline{\epsilon}, \bar{\epsilon} > 0$, $\underline{q} \in (0, 1)$, and $R' > R^*$ such that the following hold:*

- (a) *The mechanism that lets the buyer choose either good g with price p^{*g} , or the bundle of both goods at price $p^{*1} + p^{*2} - \underline{\epsilon}$, earns expected profit at least R^* for any negatively affiliated joint distribution $\pi \in \Pi$.*
- (b) *Consider the following mechanism: first randomly choose $\epsilon = \underline{\epsilon}$ with probability \underline{q} or $\epsilon = \bar{\epsilon}$ with probability $\bar{q} = 1 - \underline{q}$; then offer each good g with price p^{*g} , or the bundle of both goods at price $p^{*1} + p^{*2} - \epsilon$. This mechanism earns expected profit at least R' , for any negatively affiliated joint distribution $\pi \in \Pi$.*

The argument is a straightforward extension of that given by McAfee, McMillan, and Whinston [30] for the independent case, obtained by looking at the first-order condition for ϵ . The proof is in Appendix A.

5.2 Open sets of distributions

Another approach to sensitivity analysis, which avoids making any particular choices of moment restrictions, is to look at the set of distributions π for which separate pricing is optimal, and ask how small this set is. If it contained only the specific distribution constructed in Section 4, then our result would indeed be a knife-edge result. However, we will show here that the situation is not so extreme: in particular, when the type space is *finite*, there exists an open set of such worst-case π 's (in the natural topology on Π).

We could also simply try to show that there is an open set of possible moment restrictions z in (5.1) for which the worst-case optimal revenue is R^* . Note however that this is a much weaker statement than existence of an open set of π 's as above. In fact, as long as there is more than one π — say π_1 and π_2 — for which optimal revenue is R^* , then for any z in the open set satisfying $E_{\pi_1}[z(\theta)] < 0$ and $E_{\pi_2}[z(\theta)] > 0$, there is some convex combination of π_1 and π_2 for which $E[z(\theta)] = 0$, and so the worst-case objective (5.3) equals R^* .

To get to our non-knife-edge result, the important step is the following:

Lemma 5.2. *Consider the benchmark monopoly application. Suppose each set of values Θ^g is finite, that each π^g has full support, and also suppose that for each single-good*

problem π^g , the optimal price is unique. Then, there exists a joint distribution $\pi \in \Pi$ for which selling each good separately constitutes the unique optimal mechanism.

The proof involves a more careful retracing of the generalized-virtual-value maximization arguments in Section 4. We leave the details to Appendix A.

Now we infer the result:

Corollary 5.3. *In the setting of Lemma 5.2, there exists a set of joint distributions $\widehat{\Pi} \subseteq \Pi$, which is open (in the relative topology on Π , as a subset of $\mathbb{R}^{|\Theta|}$), and such that for any $\pi \in \widehat{\Pi}$, no joint screening mechanism earns an expected profit higher than R^* .*

Proof. As argued in the proof of Lemma 4.3 in the appendix, when looking for optimal mechanisms, we can restrict to ones whose payments are all in $[-\bar{t}, \bar{t}]$ for some sufficiently high constant \bar{t} . Then, the effective space of mechanisms \mathcal{M}' becomes a convex polytope, since it is a compact set of $|\Theta| \cdot (G + 1)$ -dimensional vectors defined by certain linear constraints. Therefore, it is the convex hull of its vertices (see e.g. [46, Theorem 1.1]), i.e. there exist some mechanisms M_1, \dots, M_K such that every mechanism in \mathcal{M}' equals some convex combination of them.

By Lemma 5.2, there exists some particular $\pi^* \in \Pi$ for which the separate-sales mechanism earns *strictly* higher expected profit than any other mechanism. Since expected profit is a linear function on \mathcal{M}' , it is maximized at one of the corners, so the separate-sales mechanism must be one of these corners, say M_1 . By continuity, for any sufficiently nearby π , M_1 still gives strictly higher expected profit than M_2, \dots, M_K , and so remains higher than any convex combination, i.e. no mechanism attains higher profit than R^* . □

We note that this result does not extend to the continuous-type case. If, for example, each Θ^g is an interval in \mathbb{R}^+ , with a unique optimal price p^{*g} that is in the interior, then for any joint distribution such that the optimal profit is R^* , it can be perturbed by an arbitrarily small amount so that offering the bundle of goods 1 and 2 at a price $p^{*1} + p^{*2} - \epsilon$ (in addition to offering each good g separately at price p^{*g}) earns strictly higher profit for small ϵ . As in [30], this can be seen by looking at the first-order condition with respect to ϵ .

However, we do still have an open set of moment restrictions z for which the worst-case optimum remains R^* , as long as there is more than one joint distribution π that pins profit down to R^* (which will be true in general).

6 Concluding comments

We conclude by briefly discussing some aspirations for future work. As described in the introduction, the main purpose of this paper has been to advance a possible new modeling approach to attack multidimensional screening problems, where the traditional approach has often given intractable, or at least complicated, models. Specifically, we have formulated a model in which assuming no interactions between the components leads to no interactions in the solution.

This is a first step. There are two natural goals for further work to build on this model and proof:

- First, to connect formally with the simple intuition sketched back in the introduction: the profit from screening separately on each component g is known independently of the joint distribution, whereas any mechanism in which the components interact seems sensitive to the joint distribution. Can our proof (or perhaps some completely different proof technique) be expressed in a way that centers on this idea?

This might seem an obvious approach, but there is some subtlety. This proof technique would consist of two basic steps: the profit from separate screening does not depend on the joint distribution; and, this independence makes it a solution to the maxmin problem. It is clear that the first of these steps depends on the assumption of additively separable preferences. But in fact, the second step depends on this assumption as well. That is: one can construct simple examples of multidimensional screening problems with the same objective (maxmin profit over joint distributions consistent with known marginals), but non-additively-separable preferences, in which there is a unique optimal mechanism, and its profit *does* vary depending on the joint distribution. (We omit the details here.) In view of this, it is unclear how one would go about proving the second step, without constructing an explicit distribution as we have done here.

- Second, to have fruitful applications of the method to new economic questions. More specifically, this would mean identifying some applied-theory question involving multidimensional screening, where a traditional model would be intractable, and where using the robust approach to obtain a clean formal solution turns out to be helpful in thinking through the economics of the problem. As noted before, the important insights in multidimensional screening are likely to arise when there are systematic

interactions between the different dimensions, which is exactly what the model in this paper has ruled out. However, this paper may provide a useful starting point by showing how to first eliminate interactions besides the one that is of interest.

Earlier work by this author followed this course of development in a moral hazard setting, taking a simple robust approach to principal-agent contracting [12] and applying it to gain traction on the more complex problem of incentivizing information acquisition [11]. To do the same in a multidimensional screening context is still a task that lies ahead.

A Omitted proofs

Proof of Proposition 3.1. For notational convenience, let us parameterize types directly by z , and write $\theta^g(z)$ for the corresponding value for good g . Thus, for any mechanism (x, t) , $x(z) \in [0, 1]^G$ indicates type z 's probability of receiving each good, and $t(z) \in \mathbb{R}$ indicates z 's payment; incentive-compatibility (2.1) becomes

$$\sum_g \theta^g(z)x^g(z) - t(z) \geq \sum_g \theta^g(z)x^g(\hat{z}) - t(\hat{z}) \quad \text{for all } z, \hat{z} \in [0, 1], \quad (\text{A.1})$$

and individual rationality (2.2) becomes

$$\sum_g \theta^g x^g(z) - t(z) \geq 0. \quad (\text{A.2})$$

The claim is that for any such mechanism, expected revenue satisfies

$$\int_0^1 t(z) dz \leq R^*.$$

Following the standard method, let $U(z) = \sum_g \theta^g(z)x^g(z) - t(z)$ denote the payoff of type z . So for $z' \geq z$, (A.1) gives

$$U(z') \geq U(z) + \sum_g (\theta^g(z') - \theta^g(z))x^g(z) \geq U(z). \quad (\text{A.3})$$

Thus, U is weakly increasing, hence differentiable almost everywhere, and equal to the integral of its derivative. Moreover, at each point of differentiability, the envelope theorem

applied to (A.3) gives us

$$\frac{dU}{dz} = \sum_g \frac{d\theta^g(z)}{dz} x^g(z) = \sum_g \frac{1}{f^g(\theta^g(z))} x^g(z).$$

Therefore,

$$U(z) = U(0) + \int_0^z \frac{dU(\tilde{z})}{d\tilde{z}} d\tilde{z} = U(0) + \sum_g \int_0^z \frac{x^g(\tilde{z})}{f^g(\theta^g(\tilde{z}))} d\tilde{z}.$$

Consequently, profit is

$$\begin{aligned} \int_0^1 t(z) dz &= \int_0^1 \left(\sum_g \theta^g(z) x^g(z) - U(z) \right) dz \\ &= \int_0^1 \left(\sum_g \theta^g(z) x^g(z) - U(0) - \sum_g \int_0^z \frac{x^g(\tilde{z})}{f^g(\theta^g(\tilde{z}))} d\tilde{z} \right) dz \\ &\leq \sum_g \left[\int_0^1 \theta^g(z) x^g(z) dz - \int_0^1 \left(\int_0^z \frac{x^g(\tilde{z})}{f^g(\theta^g(\tilde{z}))} d\tilde{z} \right) dz \right] \end{aligned}$$

where the last inequality holds because $U(0) \geq 0$ by (A.2). Switching the variables in the second integral and changing the order of integration gives

$$= \sum_g \left[\int_0^1 \left(\theta^g(z) - \frac{1-z}{f^g(\theta^g(z))} \right) x^g(z) dz \right]. \quad (\text{A.4})$$

Now, for each g , we can see that the expression in parentheses is exactly the virtual value v^g corresponding to marginal type $\theta^g(z)$. Hence, by assumption, this quantity is negative for $\theta^g(z) < \theta^{*g}$ and positive for $\theta^g(z) > \theta^{*g}$ — which correspond to $z < F^g(\theta^{*g})$ and $z > F^g(\theta^{*g})$, respectively. In particular, an upper bound for the value of (A.4) is found by taking $x^g(z)$ to be as small as possible, namely 0, when $z < F^g(\theta^{*g})$, and as large as possible, namely 1, when $z > F^g(\theta^{*g})$. Thus, we see that the principal's profit is at most

$$\begin{aligned} \sum_g \left[\int_{\theta^{*g}}^1 \left(\theta^g(z) - \frac{1-z}{f^g(\theta^g(z))} \right) dz \right] &= \sum_g \left[\int_{\theta^{*g}}^1 \frac{d}{dz} (-\theta^g(z)(1-z)) dz \right] \\ &= \sum_g \theta^{*g} (1 - F^g(\theta^{*g})). \end{aligned}$$

This is exactly the profit attained by the mechanism that sells the goods separately, with

a price of θ^{*g} for good g ; therefore it cannot exceed R^* . This establishes that the profit from any mechanism is at most R^* , as claimed. \square

Proof of Lemma 4.2. Let $\Delta = \max_{x,\theta} u(x,\theta) - \min_{x,\theta} u(x,\theta)$, and notice that for any mechanism satisfying IC, any two types' payments can differ by at most Δ . Consequently, when looking for optimal mechanisms, we can restrict attention to mechanisms whose payments are in $[-\bar{t}, \bar{t}]$, for a sufficiently high value of \bar{t} (which is independent of the distribution π): each type's payment is bounded below, because a sufficiently low payment for one type would imply low payments for all types, and then the mechanism would not be optimal (it could be improved by increasing all types' payments by a small constant without violating IR); and each type's payment is bounded above, simply by IR.

Then, the maximum expected profit, as a function of π , is the upper envelope of a uniformly bounded family of linear functions (one such function for each possible mechanism $M \in \mathcal{M}$ satisfying the bounds on payments). It readily follows that this upper envelope is Lipschitz in π , and in particular is continuous. \square

Proof of Lemma 4.3. (Adapted from [28])

It is easy to see that statement (b) of the lemma follows from (a), by integrating over each S_k in the partition; so it suffices to prove (a).

As in the proof of Lemma 4.2, we can take $\Delta = \max_{x,\theta} u(x,\theta) - \min_{x,\theta} u(x,\theta)$, and then in any mechanism, any two types' payments can differ by at most Δ . Also, put $\tau = \min\{\epsilon/6\Delta, 1\}$.

By Lipschitz continuity, there exists δ such that, whenever θ, θ' are two types with $d(\theta, \theta') < \delta$, then $|u(x, \theta) - u(x, \theta')| < \tau\epsilon/6$ for all x . We show this δ has the desired property.

Let (x, t) be any given mechanism. Let $\underline{t} = \min_{\theta} t(\theta)$. Let $S \subseteq \Delta(X) \times \mathbb{R}$ be the set of values $(x(\theta), \tau\underline{t} + (1 - \tau)t(\theta))$ for $\theta \in \Theta$, and let \bar{S} be its closure, which is compact (by the above observation on payments). Then define (\tilde{x}, \tilde{t}) by simply assigning to each type $\theta \in \Theta$ the outcome in \bar{S} that maximizes its payoff, $Eu(x, \theta) - t$. This exists by compactness. This (\tilde{x}, \tilde{t}) is a mechanism: IC is satisfied by definition, and IR is satisfied since the payments have only been reduced relative to those in (x, t) , so each type θ has the option of getting allocation $x(\theta)$ for a payment of less than $t(\theta)$, which gives nonnegative payoff.

Now, let $d(\theta, \theta') < \delta$. We know that the outcome chosen by θ' in the new mechanism can be approximated arbitrarily closely by an element of S corresponding to some type θ'' ; in particular, there exists θ'' such that

$$|Eu(\tilde{x}(\theta'), \theta) - Eu(x(\theta''), \theta)| < \frac{\tau\epsilon}{6} \quad \text{and} \quad |\tilde{t}(\theta') - [\tau\underline{t} + (1 - \tau)t(\theta'')]| < \frac{\tau\epsilon}{6}. \quad (\text{A.5})$$

Now, we know from IC for the original mechanism

$$Eu(x(\theta), \theta) - t(\theta) \geq Eu(x(\theta''), \theta) - t(\theta''), \quad (\text{A.6})$$

and by the definition of the new mechanism (\tilde{x}, \tilde{t}) ,

$$Eu(\tilde{x}(\theta'), \theta') - \tilde{t}(\theta') \geq Eu(x(\theta), \theta') - [\tau\underline{t} + (1 - \tau)t(\theta)].$$

Using (twice) the fact that $d(\theta, \theta') < \delta$, the latter inequality turns into

$$Eu(\tilde{x}(\theta'), \theta) - \tilde{t}(\theta') \geq Eu(x(\theta), \theta) - [\tau\underline{t} + (1 - \tau)t(\theta)] - \frac{\tau\epsilon}{3}.$$

Now combining with (A.5) we get

$$Eu(x(\theta''), \theta) - [\tau\underline{t} + (1 - \tau)t(\theta'')] > Eu(x(\theta), \theta) - [\tau\underline{t} + (1 - \tau)t(\theta)] - \frac{2\tau\epsilon}{3}. \quad (\text{A.7})$$

Adding (A.6) and (A.7), and canceling common terms, we get

$$\tau t(\theta'') > \tau t(\theta) - \frac{2\tau\epsilon}{3}$$

or

$$t(\theta'') > t(\theta) - \frac{2\epsilon}{3}.$$

Hence, from (A.5),

$$\begin{aligned} \tilde{t}(\theta') &> t(\theta'') - \tau(t(\theta'') - \underline{t}) - \frac{\tau\epsilon}{6} \\ &> \left(t(\theta) - \frac{2\epsilon}{3}\right) - \tau\Delta - \frac{\tau\epsilon}{6} \\ &\geq t(\theta) - \frac{2\epsilon}{3} - \frac{\epsilon}{6} - \frac{\epsilon}{6} \end{aligned}$$

which is the desired statement (a).

□

Proof of Proposition 5.1. Write $\underline{\theta}^g = \min(\Theta^g)$ and $\bar{\theta}^g = \max(\Theta^g)$.

Consider first the pricing problem in each component separately. The profit from setting price p^g is $p^g(1 - F^g(p^g))$, where F^g is the cumulative distribution function for θ^g . Since the optimal price p^{*g} is in the interior of Θ^g , the first-order condition must hold there:

$$1 - F^g(p^{*g}) - p^{*g}f^g(p^{*g}) = 0.$$

Now, for sufficiently small $\epsilon > 0$, we must have

$$\begin{aligned} & -\epsilon(1 - F^1(p^{*1}))(1 - F^2(p^{*2})) + \\ & (p^{*1} - \epsilon)(F^1(p^{*1}) - F^1(p^{*1} - \epsilon))(1 - F^2(p^{*2})) + \\ & (p^{*2} - \epsilon)(F^2(p^{*2}) - F^2(p^{*2} - \epsilon))(1 - F^1(p^{*1})) > 0. \end{aligned} \tag{A.8}$$

Indeed, when $\epsilon = 0$, this expression equals 0, and its derivative with respect to ϵ is

$$\begin{aligned} & -(1 - F^1(p^{*1}))(1 - F^2(p^{*2})) + p^{*1}f^1(p^{*1})(1 - F^2(p^{*2})) + p^{*2}f^2(p^{*2})(1 - F^1(p^{*1})) \\ & = (1 - F^1(p^{*1}))(1 - F^2(p^{*2})) \\ & \quad - (1 - F^2(p^{*2})) [1 - F^1(p^{*1}) - p^{*1}f^1(p^{*1})] \\ & \quad - (1 - F^1(p^{*1})) [1 - F^2(p^{*2}) - p^{*2}f^2(p^{*2})] \\ & = (1 - F^1(p^{*1}))(1 - F^2(p^{*2})) \\ & > 0. \end{aligned}$$

(Here the second equality comes from the first-order condition for each p^{*g} .) Write Δ for the left-hand side of (A.8) (which depends on ϵ). Let $\underline{\epsilon}$ be a small value for which (A.8) holds.

Consider the behavior of various buyer types under the separate-price mechanism, illustrated in Figure 2. Buyer types in region A buy both goods; those in regions B and D buy only good 2, while those in regions C and E buy only good 1.

Now consider the change in expected profit when the mechanism is changed to offering either separate prices (p^{*1}, p^{*2}) or $p^{*1} + p^{*2} - \epsilon$ for the bundle. Buyers whose value for both goods g is above p^{*g} (region A in the figure) now buy the bundle, paying ϵ less than before. Buyers with $\theta^1 \in [p^{*1} - \epsilon, p^{*1})$ and $\theta^2 \geq p^{*2}$ (region B) formerly bought only good 2 but now buy the bundle, paying $p^{*1} - \epsilon$ more than before. And buyers with




Figure 2: Buyer behavior under separate sales and bundling.

$\theta^2 \in [p^{*2} - \epsilon, p^{*2})$ and $\theta^1 \geq p^{*1}$ (region C) switch to buying the bundle, paying $p^{*2} - \epsilon$ more than before. These changes constitute a lower bound on the net change in profit. (In addition, some types who formerly bought nothing now buy the bundle; we ignore them.)

Thus, writing $\pi(A), \pi(B), \pi(C)$ for the measures of these regions under joint distribution π , our change in profit is at least

$$-\epsilon\pi(A) + (p^{*1} - \epsilon)\pi(B) + (p^{*2} - \epsilon)\pi(C). \quad (\text{A.9})$$

Now, for any negatively affiliated π , we have

$$\begin{aligned} \pi(B) &\geq \frac{\pi([p^{*1} - \epsilon, p^{*1}] \times \Theta^2)}{\pi([p^{*1}, \bar{\theta}^2] \times \Theta^2)} \times \pi(A) \\ &= \frac{F^1(p^{*1}) - F^1(p^{*1} - \epsilon)}{1 - F^1(p^{*1})} \times \pi(A) \end{aligned}$$

and similarly

$$\pi(C) \geq \frac{F^2(p^{*2}) - F^2(p^{*2} - \epsilon)}{1 - F^2(p^{*2})} \times \pi(A).$$

Plugging in to (A.9), our change in profit in going to the bundled mechanism is at

least

$$\begin{aligned} \pi(A) \times & \left[-\epsilon + (p^{*1} - \epsilon) \frac{F^1(p^{*1}) - F^1(p^{*1} - \epsilon)}{1 - F^1(p^{*1})} + (p^{*2} - \epsilon) \frac{F^2(p^{*2}) - F^2(p^{*2} - \epsilon)}{1 - F^2(p^{*2})} \right] \\ & = \pi(A) \times \frac{\Delta}{(1 - F^1(p^{*1}))(1 - F^2(p^{*2}))}, \end{aligned}$$

with Δ given by the left-hand side of (A.8).

Take $\epsilon = \underline{\epsilon}$ in the bundled mechanism. Recall that in this case, the corresponding value of Δ was positive; call this value $\underline{\Delta}$. This shows that the bundling mechanism, with price $p^{*1} + p^{*2} - \underline{\epsilon}$ for the bundle, earns expected profit at least as high as the R^* from separate sales, proving part (a).

All of this basically follows [30] (who considered the independent case). In our case, we obtain only a weak improvement from this bundling, because with negative affiliation, $\pi(A)$ may be zero, or arbitrarily close. This is why we must randomize the bundle price to obtain a strict improvement; we now detail this adjustment.

Take $\bar{\epsilon} = \min\{p^{*1} - \underline{\theta}^1, p^{*2} - \underline{\theta}^2\}$; without loss of generality, $\bar{\epsilon} = p^{*2} - \underline{\theta}^2$. Then, in the mechanism with bundle price $p^{*1} + p^{*2} - \bar{\epsilon}$, region E in Figure 2 disappears, and regions A and C constitute all of the area to the right of the line $\theta^1 = p^{*1}$, implying

$$\pi(C) = (1 - F^1(p^{*1})) - \pi(A).$$

Therefore, expression (A.9) is at least

$$\begin{aligned} -\bar{\epsilon}\pi(A) + (p^{*2} - \bar{\epsilon})[(1 - F^1(p^{*1})) - \pi(A)] \\ = \underline{\theta}^2(1 - F^1(p^{*1})) - p^{*2}\pi(A). \end{aligned}$$

Consequently, if ϵ is chosen to equal $\underline{\epsilon}$ with probability \underline{q} and $\bar{\epsilon}$ with probability $1 - \underline{q}$, the expected gain in profit relative to separate prices is at least

$$\begin{aligned} \underline{q} \frac{\underline{\Delta}}{(1 - F^1(p^{*1}))(1 - F^2(p^{*2}))} \pi(A) + (1 - \underline{q}) (\underline{\theta}^2(1 - F^1(p^{*1})) - p^{*2}\pi(A)) \\ = (1 - \underline{q}) \underline{\theta}^2(1 - F^1(p^{*1})) + \pi(A) \left[\underline{q} \frac{\underline{\Delta}}{(1 - F^1(p^{*1}))(1 - F^2(p^{*2}))} - (1 - \underline{q}) p^{*2} \right]. \end{aligned}$$

Evidently, if \underline{q} is chosen close enough to 1, the expression in brackets on the right will be positive. Then, for any negatively affiliated distribution π , the profit from the randomized bundling mechanism will be at least $R^* + (1 - \underline{q}) \underline{\theta}^2(1 - F^1(p^{*1}))$, which is strictly above

R^* , proving part (b). □

Proof of Lemma 5.2. We first prove the lemma in the special case where each Θ^g consists of at most two values. Write them as $\Theta^g = \{\theta_1^g, \theta_2^g\}$ with $\theta_1^g < \theta_2^g$ (or $\{\theta_1^g\}$ if there is just one value), and write (x^*, t^*) for the optimal mechanism that sells each good separately.

In the standard analysis of the one-good problem (as detailed in Appendix B), the virtual values associated to these two types are

$$\theta_1^g - \frac{\pi^g(\theta_2^g)}{\pi^g(\theta_1^g)} \times (\theta_2^g - \theta_1^g) \quad \text{and} \quad \theta_2^g.$$

If the former virtual value is negative, then the optimal price for the single good g is θ_2^g ; if it is positive, the optimal price is θ_1^g . The virtual value cannot be zero, because there would then not be a unique optimal price, contrary to assumption. In either case, the constraints that receive positive weight in the dual problem are the IR constraint for θ_1^g and the IC constraint for θ_2^g to imitate θ_1^g . (If there is just one type θ_1^g , then its virtual value is the positive number θ_1^g , and its IR constraint is binding.)

The distribution π constructed in the proof of Lemma 4.4 has full support on Θ . To see this, we look explicitly at the Markov process by which π is defined. This process consists of Poisson transitions from the low value θ_1^g to the high value θ_2^g , occurring independently across g , punctuated by Poisson resets to θ_1^g for all g simultaneously. (Again, this is for g with $|\Theta^g| = 2$; of course if $|\Theta^g| = 1$ then θ^g never moves.) It is easy to see that the stationary distribution, π , has full support. This also implies that all of the multipliers on the adjacent downward constraints constructed in the proof of Theorem 2.1, that is to say all of the $\lambda[(\theta_2^g, \theta^{-g}) \rightarrow (\theta_1^g, \theta^{-g})]$, are strictly positive, in view of their definition (4.16). Likewise, the lowest type $\theta_1 = (\theta_1^1, \dots, \theta_1^G)$ has a strictly positive weight on its IR constraint (namely $\kappa[\theta_1] = 1$).

Now, we saw in the proof of Theorem 2.1 that any mechanism (x, t) must satisfy

$$\sum_{\theta} \pi(\theta) t(\theta) \leq \sum_{\theta} \sum_g \pi(\theta) E \bar{u}^g(x^g(\theta), \theta^g) \leq \sum_{\theta} \sum_g \pi(\theta) \bar{u}_{\max}^g(\theta^g) = R^* \quad (\text{A.10})$$

with equality for the mechanism (x^*, t^*) . Since the $\pi(\theta)$ are all strictly positive, a mechanism can be optimal only if it satisfies $E \bar{u}^g(x^g(\theta), \theta^g) = \bar{u}_{\max}^g(\theta^g)$ for every θ and every g , i.e. it maximizes the virtual value in every component for every type. Since the virtual value from allocating good g is always positive or negative, never 0 (while the virtual

value from not allocating is 0), the virtual value maximizer is unique. Thus we must have $x(\theta) = x^*(\theta)$ for every θ . Moreover, we can see from (4.20) that the first inequality in (A.10) can be an equality only if all the IC and IR constraints that have positive multipliers hold with equality. Given that x and x^* coincide, this equality uniquely determines the payment of the lowest type $t(\theta_1)$, by its IR, and then uniquely determines the payment of each other type, by upward induction using the IC's. Hence we have $t(\theta) = t^*(\theta)$ for every θ as well.

This shows that the only possible optimal mechanism is $(x, t) = (x^*, t^*)$, proving the lemma in the case where each Θ^g has at most two elements.

Now we can prove the lemma in general. For each good g , write $\Theta^g = \{\theta_1^g, \dots, \theta_{j_g}^g\}$, with the values listed in increasing order, $\theta_1^g < \dots < \theta_{j_g}^g$. By assumption, the optimal price to sell good g is unique, and clearly it must equal one of the values θ_j^g ; write j^{*g} for the index, so that the optimal price is $\theta_{j^{*g}}^g$. Write (x^*, t^*) for the mechanism that sells each good g separately at price $\theta_{j^{*g}}^g$.

Let \mathcal{S}^g be the collection of all subsets of Θ^g that contain $\theta_{j^{*g}}^g$. For each such subset $\tilde{\Theta}^g \in \mathcal{S}^g$, let $\pi^g[\tilde{\Theta}^g]$ be some distribution in the corresponding one-good problem whose support is $\tilde{\Theta}^g$, and for which the unique optimal mechanism is a posted price of $\theta_{j^{*g}}^g$. (This can be constructed, for example, by placing large enough probability mass on $\theta_{j^{*g}}^g$.) Now, by choosing a sufficiently small positive weight $\eta^g[\tilde{\Theta}^g]$ for each $\tilde{\Theta}^g \in \mathcal{S}^g$, we can write π^g as a convex combination of distributions

$$\pi^g = \sum_{\tilde{\Theta}^g \in \mathcal{S}^g} \eta^g[\tilde{\Theta}^g] \pi^g[\tilde{\Theta}^g] + \eta^g[\emptyset] \bar{\pi}^g \quad (\text{A.11})$$

where $\bar{\pi}^g$ is some distribution that still has full support on Θ^g , and still has the property that the unique optimal price is $\theta_{j^{*g}}^g$. For convenience, write $\bar{\mathcal{S}}^g = \mathcal{S}^g \cup \{\emptyset\}$ and $\pi^g[\emptyset] = \bar{\pi}^g$; this allows us to write more simply

$$\pi^g = \sum_{\tilde{\Theta}^g \in \bar{\mathcal{S}}^g} \eta^g[\tilde{\Theta}^g] \pi^g[\tilde{\Theta}^g]. \quad (\text{A.12})$$

Let $\bar{\mathcal{S}} = \times_g \bar{\mathcal{S}}^g$. Consider any choice of $\tilde{\Theta} = (\tilde{\Theta}^1, \dots, \tilde{\Theta}^G) \in \bar{\mathcal{S}}$. We know that for each separate good g , setting a price of $\theta_{j^{*g}}^g$ for each item is optimal against each marginal distribution $\pi^g[\tilde{\Theta}^g]$. Accordingly, let $\pi[\tilde{\Theta}]$ be the joint distribution constructed in Subsection 4.2, so that its marginals are the distributions $\pi^g[\tilde{\Theta}^g]$ and such that (x^*, t^*) is an optimal mechanism for distribution $\pi[\tilde{\Theta}]$.

Then, we can define a joint distribution π on Θ by

$$\pi = \sum_{\tilde{\Theta} \in \mathcal{S}} \left(\prod_{g=1}^G \eta^g[\tilde{\Theta}^g] \right) \pi[\tilde{\Theta}].$$

It is straightforward to check, using (A.12), that π is a distribution on Θ whose marginal on each Θ^g equals π^g ; and (x^*, t^*) is an optimal mechanism for distribution π . We claim that in fact (x^*, t^*) is the unique optimal mechanism for π .

So let (x, t) be any optimal mechanism for π ; we wish to show that it fully coincides with (x^*, t^*) . Since $E_{\pi[\tilde{\Theta}]}[t(\theta)] \leq E_{\pi[\tilde{\Theta}]}[t^*(\theta)]$ for each $\pi[\tilde{\Theta}]$, the only way t can obtain the same expected profit as t^* against π is to have equality for every $\pi[\tilde{\Theta}]$, i.e. (x, t) must be an optimal mechanism for every distribution $\pi[\tilde{\Theta}]$.

Consider in particular any $\tilde{\Theta}$ where each $\tilde{\Theta}^g$ consists of $\theta_{j^*}^g$ and at most one other type. For these sets, the special case of the lemma we have already proven shows that we must have $(x(\theta), t(\theta)) = (x^*(\theta), t^*(\theta))$ for each $\theta \in \times_g \tilde{\Theta}^g$.

But every type $\theta \in \Theta$ belongs to some such subspace of types, for appropriate $\tilde{\Theta}$. Therefore, the optimal mechanism (x, t) must coincide with (x^*, t^*) everywhere. □

B Generalized virtual values in the monopoly problem

In this appendix we demonstrate in detail how the generalized virtual values we have defined in Subsection 4.1 reduce to the traditional ironed virtual value in the one-dimensional case. We focus on the benchmark monopoly problem with a single good (and a finite set of types). We could allow for a convex cost of production and the calculations would be virtually identical, but for simplicity we do not do so here.

Suppose that the set of possible values for the good is $\{\theta_1, \dots, \theta_J\}$, with $0 \leq \theta_1 < \theta_2 < \dots < \theta_J$, and π is the distribution. Recall the notation for allocations: $X = \{0, 1\}$, and $u(x, \theta) = x\theta$. Also write $R_j = \theta_j \times \sum_{j'=j}^J \pi(\theta_{j'})$, the profit from setting a price of θ_j ; write $R^* = \max_j R_j$, with j^* as the index attaining the maximum (if there are several, pick the lowest). It will also be convenient to put $R_{J+1} = 0$.

The traditional analysis of the problem begins by defining the virtual value of type θ_j

as

$$\tilde{\theta}_j = \theta_j - \frac{\sum_{j' > j} \pi(\theta_{j'})}{\pi(\theta_j)} (\theta_{j+1} - \theta_j), \quad \text{or } \tilde{\theta}_j = \theta_j \text{ if } j = J.$$

(This is the discrete-type analogue of the classical formula from [32]; see e.g. [44, p. 118].)

Consider first the no-ironing case, where $\tilde{\theta}_1 \leq \dots \leq \tilde{\theta}_J$. In this case, the traditional solution uses just the IR constraint of the lowest type θ_1 and the adjacent downward IC constraints $\theta_j \rightarrow \theta_{j-1}$ to show that no mechanism can earn more than R^* , which is achieved by allocating the goods to the types θ_j with nonnegative virtual value.

This method corresponds to a solution to the dual problem that puts positive weights only on these constraints. To fully illustrate the connection, we will explicitly write out what this dual solution is, and then check that our generalized virtual values defined by (4.9) correspond to the virtual values $\tilde{\theta}_j$. Recall that in a typical screening problem, there may be many possible generalized virtual values, depending on the choice of dual solution. However, once we have decided to use a dual solution that puts weight only on the lowest IR and the adjacent downward IC constraints, an easy induction using (4.7) shows that these weights are uniquely determined. Thus it makes sense to talk about *the* generalized virtual values representing this approach to the screening problem.

In our proposed dual solution, the IR and IC multipliers are as follows:

$$\kappa[\theta_1] = 1, \quad \lambda[\theta_j \rightarrow \theta_{j-1}] = \sum_{j'=j}^J \pi(\theta_{j'}),$$

and all other λ, κ variables are zero. Also, for each θ_j , we define

$$\mu_0[\theta_j] = \max\{\pi(\theta_j)\tilde{\theta}_j, 0\}, \quad \mu_1[\theta_j] = \max\{0, -\pi(\theta_j)\tilde{\theta}_j\}, \quad \nu[\theta_j] = -\max\{\pi(\theta_j)\tilde{\theta}_j, 0\}.$$

Let us check that this is indeed an optimal solution. It is immediate that all the λ, κ, μ variables are nonnegative. It is also immediate that (4.6) holds with $a = 0$ (this means not allocating the good) since all the $u(a, \theta)$ terms are zero. For $a = 1$ (allocating the good), the first three terms in (4.6) add up to $-\pi(\theta_j)\tilde{\theta}_j$, while the last two add up to $-(\max\{0, -\pi(\theta_j)\tilde{\theta}_j\} - \max\{\pi(\theta_j)\tilde{\theta}_j, 0\}) = \pi(\theta_j)\tilde{\theta}_j$. So (4.6) is satisfied. It is straightforward to check that (4.7) is satisfied as well.

And for (4.8), note that

$$\pi(\theta_j)\tilde{\theta}_j = R_j - R_{j+1}$$

for each $j = 1, \dots, J$. Since j^* is the (lowest) index for which R_j attains the maximum,

this implies $\tilde{\theta}_{j^*-1} < 0$ (if $j^* > 1$) and $\tilde{\theta}_{j^*} \geq 0$. So, since the $\tilde{\theta}_j$ are increasing, we have $\tilde{\theta}_j \geq 0$ precisely when $j \geq j^*$. Consequently, we have

$$\sum_{j=1}^J \nu[\theta_j] = - \sum_{j \geq j^*} \pi(\theta_j) \tilde{\theta}_j = - \sum_{j \geq j^*} (R_j - R_{j+1}) = -R^*.$$

So (4.8) is satisfied, and we do indeed have an optimal solution to the dual program.

Now consider the generalized virtual value of type θ_j as defined in (4.9) with these dual variables. Certainly $\bar{u}(a, \theta_j) = 0$ when $a = 0$ (all the $u(a, \theta)$ terms are zero); the relevant case is $a = 1$. For each $\theta = \theta_j$ we only have one nonzero term $\lambda[\hat{\theta} \rightarrow \theta]$, namely $\hat{\theta} = \theta_{j+1}$ (if $j = J$ there are no such terms), and then it is clear that $\bar{u}(1, \theta_j)$ is indeed equal to the traditional virtual value $\tilde{\theta}_j$.

Now we turn to the general case, where the $\tilde{\theta}_j$ are not necessarily increasing so there will be ironing. We follow the ironing procedure in [32] (with adjustments for discrete types). Consider the set of points in the plane,

$$S = \left\{ \left(\sum_{j' < j} \pi(\theta_{j'}), R_j \right) \mid j = 1, \dots, J+1 \right\}$$

and define the function $G : [0, 1] \rightarrow \mathbb{R}$ to be the upper boundary of the convex hull of this S , i.e. the lowest concave function such that $G(x) \geq y$ for each point $(x, y) \in S$. Then define the ironed revenue

$$\bar{R}_j = G \left(\sum_{j' < j} \pi(\theta_{j'}) \right)$$

for each $j = 1, \dots, J+1$. We immediately have $R_j \leq \bar{R}_j \leq \max_{(x,y) \in S} y = R^*$, for all j , with equality for $j = j^*$. Define also

$$\bar{\theta}_j = \frac{\bar{R}_j - \bar{R}_{j+1}}{\pi(\theta_j)}$$

for each $j = 1, \dots, J$. These are the ironed virtual values. Concavity of G implies they are increasing, $\bar{\theta}_1 \leq \dots \leq \bar{\theta}_J$.

Also, as in the no-ironing case, $\bar{R}_{j^*} = \max_j \bar{R}_j$ implies that $\bar{\theta}_{j^*-1} \leq 0$ (if $j^* > 1$) and $\bar{\theta}_{j^*} \geq 0$.

Now we describe how to translate the ironing approach into a solution to the dual problem. The usual colloquial description of ironing is that it maximizes revenue sub-

ject only to the adjacent downward IC constraints (and IR of the lowest type) plus a monotonicity constraint, $x(\theta_j)$ increasing in j . However, monotonicity is itself obtained as a consequence of the adjacent downward and upward IC constraints. So in fact the corresponding dual solution will put weight on both the downward and upward adjacent IC's.

For the IR constraints, we put $\kappa[\theta_1] = 1$ and $\kappa[\theta_j] = 0$ otherwise, as before. For the IC constraints, we put the following weights on the adjacent incentive constraints:

$$\begin{aligned}\lambda[\theta_j \rightarrow \theta_{j-1}] &= \sum_{j'=j}^J \pi(\theta_{j'}) + \frac{\bar{R}_j - R_j}{\theta_j - \theta_{j-1}}, \\ \lambda[\theta_{j-1} \rightarrow \theta_j] &= \frac{\bar{R}_j - R_j}{\theta_j - \theta_{j-1}}\end{aligned}$$

(and all other $\lambda[\dots]$ multipliers equal to zero). Notice that these weights are nonnegative, since $\bar{R}_j \geq R_j$. Put also

$$\mu_0[\theta_j] = \max\{\pi(\theta_j)\bar{\theta}_j, 0\}, \quad \mu_1[\theta_j] = \max\{0, -\pi(\theta_j)\bar{\theta}_j\}, \quad \nu[\theta_j] = -\max\{\pi(\theta_j)\bar{\theta}_j, 0\}.$$

In particular, the μ 's are nonnegative as well.

Let's check that this is an optimal dual solution. First let's check (4.6) (in the case $a = 1$, since $a = 0$ is easy). We'll do the case $1 < j < J$ here. Then, the first three terms on the left side of (4.6) are

$$\begin{aligned}& \lambda[\theta_{j+1} \rightarrow \theta_j] \theta_{j+1} + \lambda[\theta_{j-1} \rightarrow \theta_j] \theta_{j-1} - (\lambda[\theta_j \rightarrow \theta_{j+1}] + \lambda[\theta_j \rightarrow \theta_{j-1}]) \theta_j \\ &= \sum_{j'=j+1}^J \pi(\theta_{j'})\theta_{j+1} - \sum_{j'=j}^J \pi(\theta_{j'})\theta_j + \frac{\bar{R}_{j+1} - R_{j+1}}{\theta_{j+1} - \theta_j}\theta_{j+1} + \frac{\bar{R}_j - R_j}{\theta_j - \theta_{j-1}}\theta_{j-1} \\ &\quad - \frac{\bar{R}_{j+1} - R_{j+1}}{\theta_{j+1} - \theta_j}\theta_j - \frac{\bar{R}_j - R_j}{\theta_j - \theta_{j-1}}\theta_j \\ &= (R_{j+1} - R_j) + \frac{\bar{R}_{j+1} - R_{j+1}}{\theta_{j+1} - \theta_j}(\theta_{j+1} - \theta_j) - \frac{\bar{R}_j - R_j}{\theta_j - \theta_{j-1}}(\theta_j - \theta_{j-1}) \\ &= (R_{j+1} - R_j) + (\bar{R}_{j+1} - R_{j+1}) - (\bar{R}_j - R_j) \\ &= \bar{R}_{j+1} - \bar{R}_j \\ &= -\pi(\theta_j)\bar{\theta}_j.\end{aligned}$$

The other two terms on the left side of (4.6) equal $-(\mu_1[\theta_j] + \nu[\theta_j]) = \pi(\theta_j)\bar{\theta}_j$. Thus, (4.6) is satisfied. (This is the case $1 < j < J$, but the remaining cases are similar, using the

identities $\bar{R}_1 = R_1$, $\bar{R}_{J+1} = R_{J+1}$ to account for the missing terms.)

It is straightforward to see that (4.7) is satisfied as well: the only difference from the no-ironing case is the addition of the $(\bar{R}_{j+1} - R_{j+1})/(\theta_{j+1} - \theta_j)$ and $(\bar{R}_j - R_j)/(\theta_j - \theta_{j-1})$ terms, which each appear once with a + sign and once with a - sign on the left side, and so they cancel out.

And because $\bar{\theta}_1 \leq \dots \leq \bar{\theta}_{j^*-1} \leq 0 \leq \bar{\theta}_{j^*} \leq \dots \leq \bar{\theta}_J$, we have $\nu[\theta_j] = 0$ for $j < j^*$ and $= -\pi(\theta_j)\bar{\theta}_j$ for $j \geq j^*$. Therefore,

$$\sum_{j=1}^J \nu[\theta_j] = - \sum_{j \geq j^*} \pi(\theta_j)\bar{\theta}_j = - \sum_{j \geq j^*} (\bar{R}_j - \bar{R}_{j+1}) = -\bar{R}_{j^*} = -R^*.$$

Thus (4.8) holds as well, and we indeed have an optimal solution to the dual problem.

With this choice of dual variables, the generalized virtual value for allocating the object to type θ_j , as defined in (4.9), equals

$$\begin{aligned} \bar{u}(1, \theta_j) &= \theta_j - \left(\frac{\lambda[\theta_{j+1} \rightarrow \theta_j]}{\pi(\theta_j)} (\theta_{j+1} - \theta_j) + \frac{\lambda[\theta_{j-1} \rightarrow \theta_j]}{\pi(\theta_j)} (\theta_{j-1} - \theta_j) \right) \\ &= \theta_j - \frac{\sum_{j'=j+1}^J \pi(\theta_{j'})}{\pi(\theta_j)} (\theta_{j+1} - \theta_j) - \frac{\bar{R}_{j+1} - R_{j+1}}{\pi(\theta_j)} (\theta_{j+1} - \theta_j) - \frac{\bar{R}_j - R_j}{\pi(\theta_j)} (\theta_{j-1} - \theta_j) \\ &= \bar{\theta}_j - \frac{\bar{R}_{j+1} - R_{j+1}}{\pi(\theta_j)} + \frac{\bar{R}_j - R_j}{\pi(\theta_j)} \\ &= \frac{(R_j - R_{j+1}) - (\bar{R}_{j+1} - R_{j+1}) + (\bar{R}_j - R_j)}{\pi(\theta_j)} \\ &= \frac{\bar{R}_j - \bar{R}_{j+1}}{\pi(\theta_j)} \\ &= \bar{\theta}_j. \end{aligned}$$

(Again, this is for $1 < j < J$ but the calculation for $j = 1, J$ is almost identical.)

Thus, as promised, the generalized virtual values are equal to the ironed virtual values as traditionally defined.

References

- [1] William James Adams and Janet L. Yellen (1976), "Commodity Bundling and the Burden of Monopoly," *Quarterly Journal of Economics* 90 (3), 475–498.

- [2] Mark Armstrong (1999), “Price Discrimination by a Many-Product Firm,” *Review of Economic Studies* 66 (1), 151–168.
- [3] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg (2014), “A Simple and Approximately Optimal Mechanism for an Additive Buyer,” in *Proceedings of the IEEE 55th Symposium on Foundations of Computer Science (FOCS ’14)*, 21–30.
- [4] Yannis Bakos and Erik Brynjolfsson (1999), “Bundling Information Goods: Pricing, Profits, and Efficiency,” *Management Science* 45 (12), 1613–1630.
- [5] Erik J. Balder (1996), “On the Existence of Optimal Contract Mechanisms for Incomplete Information Principal-Agent Models,” *Journal of Economic Theory* 68 (1), 133–148.
- [6] David P. Baron and Roger B. Myerson (1982), “Regulating a Monopolist with Unknown Costs,” *Econometrica* 50 (4), 911–930.
- [7] Marco Battaglini and Rohit Lamba (2015), “Optimal Dynamic Contracting: the First-Order Approach and Beyond,” working paper, Pennsylvania State University.
- [8] Dirk Bergemann and Stephen Morris (2005), “Robust Mechanism Design,” *Econometrica* 73 (6), 1771–1813.
- [9] Benjamin Brooks (2013), “Surveying and Selling: Belief and Surplus Extraction in Auctions,” working paper, University of Chicago.
- [10] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg (2012), “Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization,” in *Proceedings of the IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS ’12)*, 130–139.
- [11] Gabriel Carroll (2012), “Robust Incentives for Information Acquisition,” working paper, Stanford University.
- [12] Gabriel Carroll (2015), “Robustness and Linear Contracts,” *American Economic Review* 105 (2), 536–563.
- [13] Chenghuan Sean Chu, Phillip Leslie, and Alan Sorensen (2011), “Bundle-Size Pricing as an Approximation to Mixed Bundling,” *American Economic Review* 101 (1), 263–303.

- [14] Kim-Sau Chung and J. C. Ely (2007), “Foundations of Dominant-Strategy Mechanisms,” *Review of Economic Studies* 74 (2), 447–476.
- [15] Jeffrey Church and Roger Ware (2000), *Industrial Organization: A Strategic Approach*. Boston: McGraw-Hill.
- [16] Vincent Conitzer and Tuomas Sandholm (2002), “Complexity of Mechanism Design,” in *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence (UAI ’02)*, 103–110.
- [17] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos (2013), “Mechanism Design via Optimal Transport,” in *Proceedings of the 14th ACM Conference on Electronic Commerce (EC ’13)*, 269–286.
- [18] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos (2014), “Complexity of Optimal Mechanism Design,” in *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA ’14)*, 1302–1318.
- [19] Emmanuel Farhi and Iván Werning (2013), “Insurance and Taxation over the Life Cycle,” *Review of Economic Studies* 80 (2), 596–635.
- [20] Sebastian Findeisen and Dominik Sachs (2015), “Redistribution and Insurance with Simple Tax Instruments,” working paper, University of Mannheim.
- [21] Alexander Frankel (2014), “Aligned Delegation,” *American Economic Review* 104 (1), 66–83.
- [22] Daniel Garrett (2014), “Robustness of Simple Menus of Contracts in Cost-Based Procurement,” *Games and Economic Behavior* 87, 631–641.
- [23] Nima Haghpanah and Jason Hartline (2015), “Multi-dimensional Virtual Values and Second-degree Price Discrimination,” working paper, Northwestern University.
- [24] Sergiu Hart and Noam Nisan (2012), “Approximate Revenue Maximization with Multiple Items,” working paper, Hebrew University of Jerusalem.
- [25] Sergiu Hart and Noam Nisan (2013), “The Menu-Size Complexity of Auctions,” working paper, Hebrew University of Jerusalem.
- [26] Sergiu Hart and Philip J. Reny (2015), “Maximal Revenue with Multiple Goods: Nonmonotonicity and Other Observations,” *Theoretical Economics* 10 (3), 893–922.

- [27] Xinye Li and Andrew Chi-Chih Yao (2013), “On Revenue Maximization for Selling Multiple Independently Distributed Items,” *Proceedings of the National Academy of Sciences* 110 (28), 11232–11237.
- [28] Kristof Madarász and Andrea Prat (2012), “Sellers with Misspecified Models,” working paper, London School of Economics.
- [29] Alejandro M. Manelli and Daniel R. Vincent (2006), “Bundling as an Optimal Selling Mechanism for a Multiple-Good Monopolist,” *Journal of Economic Theory* 127 (1), 1–35.
- [30] R. Preston McAfee, John McMillan, and Michael D. Whinston (1989), “Multiproduct Monopoly, Commodity Bundling, and Correlation of Values,” *Quarterly Journal of Economics* 104 (2), 371–383.
- [31] Michael Mussa and Sherwin Rosen (1978), “Monopoly and Product Quality,” *Journal of Economic Theory* 18 (2), 301–317.
- [32] Roger B. Myerson (1981), “Optimal Auction Design,” *Mathematics of Operations Research* 6 (1), 58–73.
- [33] John Riley and Richard Zeckhauser (1983), “Optimal Selling Strategies: When to Haggle, When to Hold Firm,” *Quarterly Journal of Economics* 98 (2), 267–289.
- [34] Jean-Charles Rochet and Lars A. Stole (2003), “The Economics of Multidimensional Screening,” in *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Vol. I*, ed. Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, Cambridge University Press.
- [35] Casey Rothschild and Florian Scheuer (2013), “Redistributive Taxation in the Roy Model,” *Quarterly Journal of Economics* 128 (2), 623–668.
- [36] Casey Rothschild and Florian Scheuer (2014), “A Theory of Income Taxation under Multidimensional Skill Heterogeneity,” working paper, Stanford University.
- [37] Michael Rothschild and Joseph Stiglitz (1976), “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *Quarterly Journal of Economics* 90 (4), 629–649.

- [38] Aviad Rubinstein and S. Matthew Weinberg, (2015), “Simple Mechanisms for a Sub-additive Buyer and Applications to Revenue Monotonicity,” in *Proceedings of the 16th ACM Conference on Electronic Commerce (EC ’15)*, 377–394.
- [39] Ilya Segal (2003), “Optimal Pricing Mechanisms with Unknown Demand,” *American Economic Review* 93 (3), 509–529.
- [40] Stefanie Stantcheva (2015), “Optimal Taxation and Human Capital Policies over the Life Cycle,” working paper, Harvard University.
- [41] George J. Stigler (1963), “United States v. Loew’s Inc.: A Note on Block-Booking,” *Supreme Court Review*, 152–157.
- [42] Roland Strausz (2006), “Deterministic Versus Stochastic Mechanisms in Principal-Agent Models,” *Journal of Economic Theory* 128 (1), 306–314.
- [43] John Thanassoulis (2004), “Haggling over Substitutes,” *Journal of Economic Theory* 117 (2), 217–245.
- [44] Rakesh Vohra (2011), *Mechanism Design: A Linear Programming Approach*, Cambridge University Press, New York.
- [45] Zihe Wang and Pingzhong Tang (2014), “Optimal Mechanisms with Simple Menus,” in *Proceedings of the Fifteenth ACM Conference on Economics and Computation (EC’14)*, 227–240.
- [46] Günter M. Ziegler (1995), *Lectures on Polytopes*, Graduate Texts in Mathematics #152, Springer-Verlag, New York.