

A CHARACTERIZATION OF “PHELPSIAN” STATISTICAL DISCRIMINATION

CHRISTOPHER P. CHAMBERS AND FEDERICO ECHENIQUE

ABSTRACT. We establish that statistical discrimination is possible if and only if it is impossible to uniquely identify the signal structure observed by an employer from a realized empirical distribution of skills. The impossibility of statistical discrimination is shown to be equivalent to the existence of a fair, skill-dependent, remuneration for workers. Finally, we connect the statistical discrimination literature to Bayesian persuasion, establishing that if discrimination is absent, then the optimal signaling problem results in a linear payoff function (as well as a kind of converse).

1. INTRODUCTION

In seminal contributions, Arrow (1971; 1973) and Phelps (1972) postulated that discrimination along racial lines, or gender identities, can have a statistical explanation. In this note we focus on Phelps’ notion of statistical discrimination: on the idea that two populations of workers, who are in essence identical, may have different economic remunerations for purely informational reasons.¹

Phelps’ theory connects worker remuneration with the distribution of signals that can be observed about worker skills. Phelps assumes a firm who observes a signal about the underlying skills of a worker. The firm observes the signal before assigning the worker to a task. The worker

(Chambers) DEPARTMENT OF ECONOMICS, GEORGETOWN UNIVERSITY
(Echenique) DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES, CALIFORNIA
INSTITUTE OF TECHNOLOGY

Echenique thanks the NSF for support through the grants SES-1558757 and CNS-1518941. We are grateful to Leeat Yariv for comments on a previous draft.

¹We follow the interpretation of Phelps’ model due to Aigner and Cain (1977). Arrow’s theory of statistical discrimination relies on a coordination failure, and is quite different from Phelps’. Statistical discrimination stands in contrast with taste-based discrimination, as in Becker (1957).

is paid her expected contribution to the firm, conditional on the firm's observed signal about the worker. (A competitive market ensures that workers are paid their contributions.) Consider now two populations of workers: group A and group B. If the signal is more informative for As than for Bs, then (the argument goes), a worker from group A may be ex-ante more valuable to the firm than a B worker. This is because the additional information about the A worker may be used to better assign her a task matching her skills. Even more, the signal may be the result of a test that has been designed with a population from group A in mind. The signal implemented by the test will then be more informative about the skills of a prospective A worker than a B worker.²

As a consequence of the difference in informativeness, the firm may value a group A worker over a group B worker. We formulate the theory of statistical discrimination using the language of the recent literature on informational design. A firm observes a signal about a worker's skills, and bases both the assigned task and the payment to the worker on the revenue it expects to gain from the action taken by the worker at the firm. A group of workers comes with a distribution over signals: an information structure. The distribution over signals of group A may be more informative than the distribution over signals of group B. We say that statistical discrimination is present if two groups of workers, each group having their distinct distribution over signals, but *the same distribution of skills*, receive different payments in expectation.

Our contribution is to connect statistical discrimination with two seemingly distinct properties of the economic environment: one is identification (in the econometric sense) of signals from skills, and the other is the linearity of firm revenue in "fair" skill-dependent payoffs. First, we show that the absence of statistical discrimination is equivalent to

²As an example, Aigner and Cain cite evidence from the education literature to the effect that the SAT is less informative about the abilities of African-American students than it is for white students.

the econometric identification of signals. Specifically, we prove that statistical discrimination is not possible if and only if every given distribution of skills arises from a unique distribution of signals. By definition, when discrimination is possible, the identification property must be violated. Our contribution is in the converse: whenever identification is impossible, discrimination can arise.

Second, we show that identification, and therefore the absence of discrimination, is equivalent to the existence of a fair skill-based remuneration for workers. Workers' payments are a linear combination of the fair remunerations. Each list of skills must be associated with a value, which is independent of any signals, and every worker is paid the expectation according to the distribution of skills inherent in her realized signal.

Finally, we show that the optimal information structure in the sense of Kamenica and Gentzkow (2011) achieves precisely the fair remuneration in our results.

2. THE MODEL

2.1. Notation. A set is *binary* if it has one or two elements. If A is a closed subset of a Euclidean space, we denote by $\Delta(A)$ the set of Borel probability measures on A .

2.2. The model. The model involves a firm and a worker. The firm faces uncertainty over the revenues it can obtain from the worker's actions. The firm's revenue depends on the worker's skills, and how those skills match up with the technology of the firm.

Let Θ denote a finite set of uncertain *states of the world*; these states represent the skill set of the worker, and are unknown to the firm. The firm asks the worker to undertake some action, and it only cares about the state-contingent payoff that results from the worker's action. Formally, then, an *action* is an element $a \in \mathbf{R}^\Theta$. Thus, the task of the firm is to properly match a worker to an action with the appropriate skill set.

There is a closed set of *signals*, or *payoff-relevant types*, $\mathcal{S} \subseteq \Delta(\Theta)$. Here we identify signals with the posterior distribution that they induce over Θ . The firm observes $s \in \mathcal{S}$ before asking the worker to undertake an action. Thus, the goal of the firm is to choose the appropriate action for the appropriate worker, after a signal of worker skill has been observed.

The firm solves the following problem. For a given $s \in \mathcal{S}$, and finite set of actions A ,

$$v_A(s) \equiv \max_{a \in A} \sum_{\theta \in \Theta} a(\theta) s(\theta).$$

Given signal $s \in \mathcal{S}$, $v_A(s)$ is the maximal expected revenue the employer can achieve. We maintain the assumption that labor markets are competitive, and therefore a worker of type s is paid the revenue $v_A(s)$ that she generates for the firm. This is as in Phelps (1972) and Aigner and Cain (1977). Observe that v_A is the “value function” of A , as in Blackwell (1953) or Machina (1984), and is thus always convex.

A probability $\pi \in \Delta(\mathcal{S})$ is an *information structure*. It induces a probability over Θ via: $p_\pi(\theta) = \int_{\mathcal{S}} s(\theta) d\pi(s)$. For a set $E \subseteq \mathcal{S}$, we can interpret $\pi(E)$ as an empirical frequency of individuals who generate signals $s \in E$. The empirical frequency π then generates an empirical frequency of skills, which is p_π .

We say that the set of signals \mathcal{S} is *non-discriminatory* if for any information structures $\pi, \pi' \in \Delta(\mathcal{S})$, and any finite set $A \subseteq \mathbf{R}^\Theta$, if $p_\pi = p_{\pi'}$, then

$$\int_{\mathcal{S}} v_A(t) d\pi(t) = \int_{\mathcal{S}} v_A(t) d\pi'(t).$$

Interpret $\pi(E)$ as the frequency of individuals of type $E \subseteq \mathcal{S}$. Under the competitive markets assumption, the set \mathcal{S} being non-discriminatory means that the average remuneration paid to a class of workers with distribution π ultimately depends only on the distribution of their skills.

2.3. Motivation and a Phelpsian example. We start by a simple example to recreate the point made by Phelps (1972). It is a minimal example; the simplest we can think of that delivers the Phelpsian message. Let $\Theta = \{\theta_1, \theta_2, \theta_3\}$ be the set of states, and $A =$

$\{(1, 0, 0), (0, 1/2, 3)\}$ be the set of available actions. Observe that with this specification, workers are not “high” or “low” quality, but they simply have differing aptitudes for the available actions.

Suppose that

$$\mathcal{S} = \{(1, 0, 0), (1/2, 1/2, 0), (0, 1/2, 1/2), (0, 0, 1)\}$$

is the set of signals, or worker types.

Consider two information structures, π and π' , described in the table below, together with the profit function v_A resulting from our assumed Θ and A :

	$t = (1, 0, 0)$	$t = (1/2, 1/2, 0)$	$t = (0, 1/2, 1/2)$	$t = (0, 0, 1)$
$\pi(t)$	1/3	0	2/3	0
$\pi'(t)$	0	2/3	0	1/3
$v_A(t)$	1	1/2	7/4	3

There are two populations of workers, say A and B. The two populations differ in the information that the firm obtains about their skills. The workers might take a test, as in Phelps (1972), and the informational content of the test might be different for the two populations. So As emit signals about their skills as given by π , while Bs distribution over signals is π' . Observe that $p_\pi = p_{\pi'} = (1/3, 1/3, 1/3)$, reflecting that the populations overall have the same skills.

A worker from group A reveals that she is either good for action $a_1 = (1, 0, 0)$ or action $a_2 = (0, 1/3, 3)$. The B worker reveals the same kind of information, but less efficiently: a signal $t = (1/2, 1/2, 0)$ tells the employer that a_1 is the optimal choice given the information at hand, but leaves the employer with some doubts as to whether a_2 may have been the optimal action. In consequence, we have

$$\int_T v_A(t) d\pi(t) = 1/3 + 7/6 > 1/3 + 1 = \int_T v_A(t) d\pi'(t).$$

If workers are paid according to the revenues that they contribute to the firm, as would be the case in a competitive market, then A workers are paid more than B workers in aggregate. The differences in expected

(or population) remuneration between the two is purely a consequence of the informational content in their corresponding signal structures.

In our example of Phelpsian statistical discrimination, the two different information structures have the same mean. This is a necessary requirement for the existence of statistical discrimination. It is important to point out, however, that skill can *always* be inferred from wages, even when there is discrimination. We present Proposition 1 to make this point.

For a set of actions $A = \{a_1, \dots, a_n\}$, and action k , let $A + k = \{a_1 + k, \dots, a_n + k\}$.

Proposition 1. *For any \mathcal{S} and any set of actions A , if $\pi, \pi' \in \Delta(\mathcal{S})$ for which $p_\pi \neq p_{\pi'}$, then there is k for which*

$$\int_T v_{A+k}(t) d\pi(t) \neq \int_T v_{A+k}(t) d\pi'(t).$$

2.4. When discrimination is impossible. Our discussion suggests that discrimination is tied to identification. Skills are always identified from payoffs, even when there is discrimination (Proposition 1). The problem is the converse identification: Here we show that the absence of discrimination is equivalent to the ability to estimate skills from signals. Importantly, we show that this can only happen when payments are linear in signals. So the absence of discrimination is equivalent to the existence of a state-dependent, signal-independent, “fair” payoff. Payments equal the expected value of such a payoff, and are called fair valuations.

We say that \mathcal{S}

- is *identified* if for any $\pi, \pi' \in \Delta(\mathcal{S})$, if $p_\pi = p_{\pi'}$, then $\pi = \pi'$;
- *admits fair valuations* if for any finite subset $A \subseteq \mathbf{R}^\Theta$, there is $\alpha_A \in \mathbf{R}^\Theta$ for which for all $t \in \mathcal{S}$,

$$v_A(t) = \sum_{\theta} \alpha_A(\theta) t(\theta).$$

- *admits fair valuations for binary sets* if for any binary subset $A \subseteq \mathbf{R}^\Theta$, there is $\alpha_A \in \mathbf{R}^\Theta$ for which for all $t \in \mathcal{S}$, $v_A(t) = \sum_{\theta} \alpha_A(\theta) t(\theta)$.

The notion that \mathcal{S} admits fair valuations captures the idea that any individual is paid according to her expected skill. Thus, for A , $\alpha_A(\theta)$ represents the value to the firm with technology A of skill set $\theta \in \Theta$, and if an individual sends signal s then she is paid the expected value of α_A according to s . Importantly, if $\pi \in \Delta(\mathcal{S})$, then

$$\int v_A(s) d\pi(s) = \alpha_A \cdot \int s d\pi(s) = \alpha_A \cdot p_\pi.$$

So, under fair valuations, the expected payment to a population of agents with information structure π only depends on the distribution of skills in that population.

Finally, say that \mathcal{S} is *non-discriminatory for binary sets* if for any $\pi, \pi' \in \Delta(\mathcal{S})$ and any binary $A \subseteq \mathbf{R}^\Theta$, if $p_\pi = p_{\pi'}$, then

$$\int_{\mathcal{S}} v_A(t) d\pi(t) = \int_{\mathcal{S}} v_A(t) d\pi'(t).$$

Theorem 2. *The following are equivalent.*

- (1) \mathcal{S} is non-discriminatory.
- (2) \mathcal{S} is non-discriminatory for binary sets.
- (3) \mathcal{S} is identified.
- (4) \mathcal{S} admits fair valuations.
- (5) \mathcal{S} admits fair valuations for binary sets.

The main import of Theorem 2 is that there is a α_A , independent of the signal s , so that the optimal contribution of the worker to the firm is the expected value of α_A . The worker is therefore remunerated according to some “fundamental” value α_A , and receives the expectation of α_A according to the signal s .

Proposition 3. *If \mathcal{S} admits fair valuations, then for each finite $A \subseteq \mathbf{R}^\Theta$ and corresponding $\alpha_A \in \mathbf{R}^\Theta$, we have for every $s^* \in \mathcal{S}$:*

$$\sum_{\theta} \alpha_A(\theta) s^*(\theta) = \inf \left\{ \sum_{\theta} y(\theta) s^*(\theta) : y \in \mathbf{R}^\Theta \text{ and } v_A(s) \leq \sum_{\theta} y(\theta) s(\theta) \forall s \in \mathcal{S} \right\}.$$

Proposition 3 means that the value of a worker with type s^* to the firm is the minimum expected payment that guarantees the worker a payoff of at least $v_A(s)$, for all signals $s \in \mathcal{S}$. This is a kind of

participation, or individual rationality, constraint. The worker may be able to guarantee a payment of $v_A(s)$ on the market, if her signal is s , and thus a firm must guarantee at least $v_A(s)$ in its choice of the “fair” payoff $\alpha_A \in \mathbf{R}^\Theta$.

2.5. Connection to Bayesian persuasion. The recent literature on Bayesian persuasion (Kamenica and Gentzkow (2011)) deals with the optimal design of information structures. It turns out that the value of optimal information design is linear if and only if \mathcal{S} admits no discrimination.

We now focus a bit more in depth on the notion of signal structure. As in Blackwell (1953), there is a natural notion of “comparative informativeness” for $\pi, \pi' \in \Delta(\mathcal{S})$. We say that π is *more informative* than π' if for every A , $\int v_A(t) d\pi(t) \geq \int v_A(t) d\pi'(t)$. Most economists will have heard of the notion of a “mean-preserving spread;” π turns out to be more informative than π' if it consists of a mean-preserving spread of π' .

We know that optimal information design will never utilize signal structures that are dominated according to the more informativeness order. As a result, optimal information structures will place probability zero on signals that can be obtained as the mean of other signals. Formally, an optimal information structure will have support on the extreme points of the convex hull of \mathcal{S} .

Now, let T be the closed convex hull of \mathcal{S} . An information structure is any probability measure $\pi \in \Delta(T)$. Then define $W_A : T \rightarrow \mathbf{R}$ via

$$W_A(s) \equiv \max\left\{ \int_T v_A(\tilde{s}) d\pi(\tilde{s}) : \pi \in \Delta(T) \text{ and } s = \int_T \tilde{s} d\pi(\tilde{s}) \right\}.$$

$W_A(s)$ is the value of an optimal information structure for a population with skill distribution s . In the following, ∂T denotes the extreme points of T ; those points which are not convex combinations of other points in T .

Return to our motivating “Phelpsian” example. There, discrimination was present even though \mathcal{S} consisted of the extreme points of its convex hull T , and thus \mathcal{S} was maximally informative. Phelps’ original

point can thus be refined: discrimination obtains because an employer has “different” information about two classes of individuals, rather than “better” information.

Let us see how this manifests itself in the choice of optimal information structure. In this case, for each $s \in \mathcal{S}$, we have (clearly) $W_A(s) = v_A(s)$, as each s is extreme in the convex hull of T . We therefore obtain: $(2/3)W_A(1/2, 1/2, 0) + (1/3)W_A(0, 0, 1) = \frac{4}{3} < \frac{3}{2} = (1/3)v_A(1, 0, 0) + (2/3)v_A(0, 1/2, 1/2) \leq W_A(1/3, 1/3, 1/3)$.

Hence, W_A is nonlinear in this case. This is a general artifact of non-identification and discrimination, as is evidenced by the following result.

Corollary 4. *For any \mathcal{S} , ∂T is non-discriminatory iff for every A , W_A is affine (linear).³*

As in Kamenica and Gentzkow (2011), W_A is always weakly concave, which admits the possibility that it is affine. Corollary 4 says that discrimination is possible exactly when W_A exhibits strict concavities.

3. CONCLUSION

We have formulated Phelps’ theory of statistical discrimination using the modern language of information design. Our results shed new light on the nature of discrimination, and on some of the empirical approaches one might take to establish the existence of statistical discrimination.

Statistical discrimination turns out to be equivalent to the absence of econometric identification of signals from skills. While the identification of skills from salaries is always possible, even in the presence of discrimination, we show that the crucial identification property is that of signals from skills.

In second place, we connect discrimination with the source of worker remunerations. We show that identification is impossible if and only if

³Because the domain of W_A is a set of probability measures, W_A is linear if it is affine. In fact, in this case we have $W_A(s) = \sum_{\theta \in \Theta} \alpha_A(\theta)s(\theta)$, where α_A is as in Proposition 3.

remunerations are linear in “fair” skill-dependent, signal-independent, payoffs.

Our results have immediate consequences for empirical research on discrimination. They imply that discrimination is absent if and only if empirical approaches to linearly estimating fair skills-based payoffs are viable.

4. PROOFS

Let T be the closed convex hull of \mathcal{S} . Recall that ∂T denotes the extreme points of T . The definition of v_A extends to T . Let $Y_A : T \rightarrow \mathbf{R}$ be the concave envelope of v_A , defined as the pointwise infimum of the affine functions that dominate v_A . So if $\mathcal{A}(T)$ denotes the space of all affine functions on T , then $v_A(t) = \inf\{l(t) : l \in \mathcal{A}(T) \text{ and } v_A \leq l\}$. Recall the definition of W_A from Section 2.5.

Lemma 5. $Y_A = W_A$

Proof. Let $l : T \rightarrow \mathbf{R}$ be an affine function and $v_A \leq l$. For any $\pi \in \Delta(T)$ with $\int_T q d\pi(q) = p$,

$$\int_T v_A(q) d\pi(q) \leq \int_T l(q) d\pi(q) = l\left(\int_T q d\pi(q)\right) = l(p),$$

as l is affine. Thus $W_A \leq l$, as π was arbitrary. This implies that $W_A \leq Y_A$, as l was arbitrary.

Now suppose that $W_A(p) < Y_A(p)$. Recall that W_A is concave. Then the set $D = \{(q, y) \in T \times \mathbf{R} : y \leq W_A(q)\}$ is closed and convex, so there exists α with $(q, y) \cdot \alpha \leq (p, W_A(p)) \cdot \alpha < (p, y') \cdot \alpha$ for all $(q, y) \in D$ and all $y' \geq Y_A(p)$. Write $\alpha = (\alpha^1, \alpha^2) \in \mathbf{R}^\Theta \times \mathbf{R}$. Clearly we cannot have $\alpha^2 = 0$ as $(p, W_A(p)) \in D$. Consider the affine function $l : T \rightarrow \mathbf{R}$ defined by

$$q \mapsto (1/\alpha^2)((p, W_A(p)) \cdot \alpha - \alpha^1 \cdot q).$$

This means that $l(p) = W_A(p) < Y_A(p)$. Moreover, for any $q \in T$, $\alpha \cdot (q, W_A(q)) \leq \alpha \cdot (p, W_A(p))$; hence,

$$l(q) = (1/\alpha^2)\alpha^1 \cdot p + W_A(p) - (1/\alpha^2)\alpha^1 \cdot q \geq W_A(q) \geq v_A(q),$$

where the last inequality follows from the definition of W_A . Then $l \in \mathcal{A}(T)$, $v_A \leq l$, and $l(p) < Y_A(p)$; a contradiction. \square

4.1. Proof of Theorem 2. By the Choquet-Meyer Theorem (Theorem II.3.7 in Alfsen (2012) or p. 56-57 in Phelps (2000)), T is a simplex iff ∂T is identified.

Now, to prove the theorem: it is obvious that $3 \implies 1 \implies 2$. We shall prove that $2 \implies 3$. To this end, let \mathcal{S} be non-discriminatory for binary menus. The proof that $2 \implies 3$ has two parts. The first is to show that $\mathcal{S} = \partial T$. The second is that T must be a simplex.

First, it is obvious by definition of T that $\partial T \subseteq \mathcal{S}$. So we prove that $\mathcal{S} \subseteq \partial T$. To this end, suppose by means of contradiction that there is $s^* \in \mathcal{S}$ for which there are $t, t' \in T$, $t \neq t'$, and $\gamma \in (0, 1)$ for which $s^* = \gamma t + (1 - \gamma)t'$. Let $f = (s^* - t) + [t \cdot s^* - s^* \cdot s^*]\mathbf{1}$ and $g = -f$. Observe that $f \cdot s^* = 0$, $g \cdot t = -t \cdot (s^* - t) - s^* \cdot (t - s^*) > 0$ and $f \cdot t' = (s^* - t) \cdot (t' - s^*) = \gamma(1 - \gamma)(t' - s^*) \cdot (t' - s^*) > 0$.

Let $A \equiv \{f, g\}$. Then we obtain that $v_A(t) \geq g \cdot t > 0$, $v_A(t') \geq f \cdot t' > 0$, while $v_A(s^*) = 0$ (as $f \cdot s^* = g \cdot s^* = 0$).

Now, for each of t, t' , there are finitely supported (by Caratheodory's theorem) π_t and $\pi_{t'}$ on ∂T (so in particular on \mathcal{S}) for which $t = \int_{\mathcal{S}} s d\pi(s)$ and $t' = \int_{\mathcal{S}} s d\pi'(s)$. This means that $\int_{\mathcal{S}} v_A(s) d\pi(s) \geq v_A(t) > 0$ and $\int_{\mathcal{S}} v_A(s) d\pi'(s) \geq v_A(t') > 0$, as v_A is convex. Then

$$\int_{\mathcal{S}} v_A(s) d(\gamma\pi + (1 - \gamma)\pi')(s) > 0.$$

But this contradicts 2 as $\int_{\mathcal{S}} s d(\gamma\pi + (1 - \gamma)\pi')(s) = \gamma t + (1 - \gamma)t' = s^*$ and $v_A(s^*) = 0$.

So we have shown that $\mathcal{S} = \partial T$, and we turn to the proof that T is a simplex (and thus \mathcal{S} identified). By Alfsen (2012) Theorem II.4.1, since T is convex and compact, T is a simplex if and only if $\mathcal{A}(T)$ forms a lattice in the usual (pointwise) ordering on functions. So, suppose by means of contradiction that $\mathcal{A}(T)$ does not form a lattice. Then, there are $f, g \in \mathcal{A}(T)$ which possess no supremum in $\mathcal{A}(T)$.

Lemma 6. *Let $f, g \in \mathcal{A}(T)$. For any $z \in \partial T$, if $f(z) \geq g(z)$, then there is $h \in \mathcal{A}(T)$ for which $h \geq f, g$ and $h(z) = f(z)$.*

Proof. Let M be the subgraph of the concave envelope of $v_{\{f,g\}}$. Observe by definition that it is the convex hull of the points $\{(z, v_{\{f,g\}}) : z \in \mathcal{S}\}$, so that it is polyhedral (Corollary 19.1.2 of Rockafellar (1970)). Therefore, by definition of polyhedral concave function, there is h supporting it at $(z, f(z))$. \square

From Lemma 6, and the fact that f and g possess no supremum in $\mathcal{A}(T)$, it follows that there is no affine function h for which for all $z \in \partial T$, $h(z) = \max\{f(z), g(z)\}$. Consequently, if $A \equiv \{f, g\}$, then Y_A is not affine, since for all $z \in \partial T$, it follows that $Y_A(z) = \max\{f(z), g(z)\} = v_A(z)$. Now, Y_A being concave and not affine means that there is $\hat{\pi} \in \Delta(T)$ with $\int_{\mathcal{S}} Y_A(q) d\hat{\pi}(q) < Y_A(p_{\hat{\pi}})$. Since $\mathcal{S} = \partial T$, and Y_A is concave, we can in fact find (by Lemma 4.1 in Phelps (2000)) $\pi \in \Delta(\mathcal{S})$ with $p_{\hat{\pi}} = p_{\pi}$ and

$$\int_{\mathcal{S}} v_A(q) d\pi(q) = \int_{\mathcal{S}} Y_A(q) d\pi(q) \leq \int_{\mathcal{S}} Y_A(q) d\hat{\pi}(q) < Y_A(p_{\hat{\pi}}) = Y_A(p_{\pi}),$$

where the first equality follows from $v_A(q) = Y_A(q)$ for $q \in \mathcal{S}$, and the second inequality from the choice of π .

Now, by Lemma 5, $Y_A(p_{\pi}) = \sup\{\int v_A(q) d\tilde{\pi}(q) : \tilde{\pi} \in \Delta(T) \text{ and } p_{\tilde{\pi}} = p_{\pi}\}$. Then there is $\pi' \in \Delta(\mathcal{S})$ (as the sup is achieved for a measure with support in $\partial T = \mathcal{S}$) with $p_{\pi} = p_{\pi'}$ and $\int_{\mathcal{S}} v_A(q) d\pi(q) < \int_{\mathcal{S}} v_A(q) d\pi'(q)$, contradicting the fact that \mathcal{S} is non-discriminatory for binary menus.

4.2. Proof of Proposition 3. The Lagrangian for the maximization problem in the definition of W_A is

$$\begin{aligned} L(\pi, \lambda) &= \int_T v_A(t) d\pi(t) + \lambda \cdot \left[p - \int_T q d\pi(q) \right] \\ &= \lambda \cdot p + \int_T (v_A(t) - \lambda \cdot p) d\pi(t) \end{aligned}$$

and apply the maximin theorem (see for example Theorem 6.2.7 in Aubin and Ekeland (2006), which applies here because $\Delta(T)$ is compact).

4.3. Proof of Proposition 1. Observe that for any A and any action l , we have $v_{A+l}(t) = v_A(t) + l \cdot t$. Now, since $p_\pi \neq p_{\pi'}$, there is l for which $l \cdot p_\pi \neq l \cdot p_{\pi'}$. Consequently, there is α for which:

$$\alpha l \cdot (p_\pi - p_{\pi'}) \neq \int_T v_A(t) d\pi'(t) - \int_T v_A(t) d\pi(t).$$

Let $k = \alpha l$, and conclude that:

$$\int_T v_{A+k}(t) d\pi(t) = k \cdot p_\pi + \int_T v_A(t) d\pi(t) \neq k \cdot p_{\pi'} + \int_T v_A(t) d\pi'(t) = \int_T v_{A+k}(t) d\pi'(t).$$

4.4. Proof of Corollary 4. By the Choquet-Meyers Theorem (Theorem II.3.7 in Alfsen (2012)) T is a simplex iff the concave envelope of every lower semicontinuous and convex function is affine. Clearly, when \mathcal{S} is identified, T is a simplex, and since v_A is convex and lower semicontinuous, we obtain that $W_A = Y_A$, the concave envelope. So W_A is affine.

Conversely, suppose that W_A is affine for each finite A . We will show that T is a simplex (so that ∂T forms the vertices of a simplex, and is identified). But this again follows from the fact that W_A is the smallest concave function on T dominating each $a \in A$. Since it is affine, it follows that $\mathcal{A}(T)$ is a lattice, and hence T is a simplex.

REFERENCES

- AIGNER, D. J. AND G. G. CAIN (1977): “Statistical theories of discrimination in labor markets,” *Industrial and Labor Relations Review*, 30, 175–187.
- ALFSEN, E. M. (2012): *Compact convex sets and boundary integrals*, vol. 57, Springer Science & Business Media.
- ARROW, K. J. (1971): “Some models of Racial Discrimination in the Labor Market,” Tech. Rep. RM-6253-RC, RAND.
- (1973): *The Theory of Discrimination. S. 3–33 in: O. Ashenfelter/A. Rees (Hrsg.), Discrimination in Labor Markets*, Princeton University Press.
- AUBIN, J.-P. AND I. EKELAND (2006): *Applied nonlinear analysis*, Courier Corporation.

- BECKER, G. S. (1957): *The Theory of Discrimination*, University of Chicago Press.
- BLACKWELL, D. (1953): “Equivalent comparisons of experiments,” *The annals of mathematical statistics*, 265–272.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian persuasion,” *American Economic Review*, 101, 2590–2615.
- MACHINA, M. J. (1984): “Temporal risk and the nature of induced preferences,” *Journal of Economic Theory*, 33, 199–231.
- PHELPS, E. S. (1972): “The Statistical Theory of Racism and Sexism,” *American Economic Review*, 62, 659–661.
- PHELPS, R. R. (2000): *Lectures on Choquet’s theorem, second edition*, Springer Science & Business Media.
- ROCKAFELLAR, R. T. (1970): *Convex analysis*, Princeton university press.