Transfer Program Complexity and the Take Up of Social Benefits[1]

Henrik Jacobsen Kleven
London School of Economics and CEPR

Wojciech Kopczuk
Columbia University, NBER and CEPR

May 2010

[1]We would like to thank the editor Alan Auerbach and two anonymous referees for detailed and constructive comments. We are also grateful to Douglas Bernheim, Jes Winther Hansen, Claus Kreiner, Joel Slemrod, and numerous seminar participants for helpful discussions and comments. Financial support from a grant from Economic Policy Research Network (EPRN) is gratefully acknowledged. Contact information for authors: (1) Henrik J. Kleven, Department of Economics, London School of Economics and Political Science (LSE), Houghton Street, London WC2A 2AE, United Kingdom. E-mail: h.j.kleven@lse.ac.uk. (2) Wojciech Kopczuk, Department of Economics, Columbia University, 1022 International Affairs Building, MC 3308, 420 W118th Street, New York, NY 10027, United States. E-mail: wojciech.kopczuk@columbia.edu.

Abstract

This paper models complexity in social programs as a byproduct of efforts to screen between deserving and undeserving applicants. While a more rigorous screening technology may have desirable effects on targeting efficiency, the associated complexity introduces transaction costs into the application process and may induce incomplete take up. The paper integrates the study of take up with the study of classification errors of type I and type II, and argues that incomplete take up can be seen as a form of type I error. We consider a government interested in ensuring a minimum income level for as many deserving individuals as possible, and characterize optimal programs when policy makers can choose the rigor of screening (and associated complexity) along with a benefit level and an eligibility criterion. It is shown that optimal program parameters reflect a trade-off at the margin between type I errors (including non-takeup) and type II errors. Optimal programs that are not universal always feature a high degree of complexity. Although it is generally possible to eliminate take up by the undeserving (type II errors), policies usually involve eligibility criteria that make them eligible and rely on complexity to restrict their participation. Even though the government is interested only in ensuring a minimum benefit level, the optimal policy may feature benefits that are higher than this target minimum. This is because benefits generically screen better than either eligibility criteria or complexity.

1 Introduction

The United States operates a large number of social programs offering support to those in need. This includes cash assistance to the poor, food stamps, health insurance, housing programs, child care support, and social security to the aged, blind and disabled. We observe several differences in the design and outcomes of these programs. One difference lies in the degree of *targeting* to selected groups of individuals viewed as 'deserving'. Although the U.S. welfare state in general relies on a much higher degree of targeting than most other countries, there is substantial variation in targeting across different programs within the U.S. At one end of the spectrum, the Medicare program is almost universal, while at the other end of the spectrum, disability insurance programs serve a relatively small population satisfying very stringent eligibility criteria. A second difference lies in the way social programs are administered and in their degree of *complexity*. Targeted programs tend to be characterized by a substantial amount of complexity and administrative hassle, whereas universal programs are simpler and more transparent. A third difference lies in the *take up* of social benefits. Incomplete take up among intended recipients is an important issue in all means-tested programs in the U.S., but there is huge variation in participation across different programs.[1]

This paper sets out a theoretical framework that facilitates an analysis of targeting, complexity and take up in social programs. We contribute to the existing literature along three dimensions. First, we take an initial step towards modeling and analyzing complexity in public programs. We go beyond viewing complexity as a negative side-effect of targeted programs, and treat it instead as a policy instrument that is chosen alongside benefit levels and eligibility rules in the design of a program. Second, we explain why governments may want to design a program with high complexity and incomplete take up by eligibles even though they have access to policy instruments which could increase take up. Third, we integrate the study of take up with the study of classification errors of type I (false rejections) and type II (false awards) in benefit award processes. In fact, we argue that non-enrollment in social programs can be seen as a form of type I error, and that it has to be understood by considering the trade-off with the usual type I and II errors.

Empirical economists have long been concerned with the issue of incomplete take-up rates in public programs. The empirical literature hypothesizes three possible explanations for incomplete take up: welfare stigma, transaction costs, and imperfect information. The seminal work in this area is the Moffitt (1983) model of welfare stigma, suggesting that eligibles may find non-participation in a welfare program optimal because it is viewed as demeaning and shameful. But the stigma hypothesis is consistent with other more concrete costs associated with taking up social benefits. Indeed, a substantial amount of evidence have documented that applying for welfare benefits involves large transaction costs arising from application processes being complex, tedious and time-consuming (Moffitt, 2003; Currie, 2006).

[1]See the survey of literature on take up by Currie (2006) and papers on specific means-tested programs in Moffitt (2003).

The complexity of welfare programs may arise from detailed eligibility criteria, rigorous documentation requirements, difficult and time-consuming forms, or requiring multiple trips to the program office for interviewing and testing. Moreover, some programs involve frequent re-certification to continue to receive the benefit, and applicants are frequently rejected because they fail to fulfill the administrative requirements within the required time. Notice that these forms of complexity reflect, at least in part, an attempt of program administrators to monitor true eligibility accurately, and hence complexity may have some desirable effects on the magnitude of classification errors. At the same time, these monitoring activities introduce hassle and possibly cognitive costs into the application process, which may hurt take up. Indeed, empirical research has shown that complexity and administrative hassle do reduce program enrollment (Currie and Grogger, 2001; Bitler et al., 2003; Daly and Burkhauser, 2003; Aizer, 2007), and that such effects may be more important than stigma (Currie, 2006).

Despite the fact that complexity and administration seem to be very important for the effects of public policies in general, and for the take up of social benefits in particular, we are not aware of theoretical work modeling the complexity of public policy. Instead, the literature on mechanism design has focused on the generosity and structure of benefits and the incentives for ineligibles to reveal themselves truthfully. The key assumption in this literature is that innate ability is unobservable at any cost, whereas earnings are perfectly observable at no cost. The government has no access to a monitoring technology to assess true eligibility and therefore has to rely on limitations in earnings-based benefits to induce self-revelation. Our paper goes beyond this extreme assumption about information by modeling the information collection process — the monitoring technology — used to elicit true eligibility for social benefits. The empirical work by Benítez-Silva et al. (2004) on disability insurance programs in the U.S. demonstrates that the monitoring technology can be a very important aspect of program design.

The way we model the complexity of a monitoring technology is consistent with the evidence discussed above. It is an instrument used by program administrators to increase the rigor of screening in order to extract a better signal of true eligibility, but which makes the application process more costly and therefore may induce non-participation by eligibles. The cost of complexity can be interpreted as a cognitive cost of having a lot of testing, or it can be interpreted as a more concrete cost from the time and money spent applying. Modeling complexity as a cost to the individual is not the only possible approach to complexity that one may consider — issues related to bounded rationality and bounded willpower may be important and at best our framework fits them only as a "reduced form" representation.

Our model also accounts for imperfect information about eligibility on part of the potential welfare applicants. The role of imperfect information for non-enrollment into social programs is well-documented (e.g. Daponte et al., 1999; Heckman and Smith, 2004), and it serves to reinforce the importance of complexity in the decision to apply for welfare. It is exactly because of imperfect information about eligibility that an individual may be reluctant to incur the transaction costs associated with applying.

The paper characterizes program parameters in equilibrium when policy makers can choose standard policy instruments — a benefit level and an eligibility rule — along with the additional instrument capturing program complexity. The model assumes that the government is interested in income maintenance, i.e. ensuring a minimum income level for as many truly poor ('deserving') individuals as possible, being constrained by a limited budget. Income maintenance is chosen as the policy objective instead of social welfare maximization, because it simplifies the analysis and sharpens the theoretical results. Furthermore, it is often argued that income maintenance is more consonant with real-world policy problems and therefore adds positive content to the analysis (Kanbur, 1987; Besley and Coate, 1992, 1995; Kanbur et al., 1994).[2] The implications of adopting a welfarist approach are discussed in Appendix B, which argues that our main qualitative findings are consistent with such an extension.

We show that optimal program parameters reflect a trade-off at the margin between type I errors (including non-takeup) and type II errors. Optimal programs that are not universal always feature a high degree of complexity. Although it is generally possible to eliminate take up by the undeserving (type II errors), policies usually involve eligibility criteria that make them eligible and rely on complexity to restrict their participation. These policies feature incomplete take up by the deserving along with classification errors of both type I and II in the benefit award process. Even though the government is interested only in ensuring a minimum benefit level, the optimal policy may feature benefits that are higher than this target minimum. This is because benefits generically screen better than either eligibility criteria or complexity.

The rest of the paper is organized as follows. Section 2 defines the different classification errors in public programs and discusses how they have been studied in the literature. Section 3 presents our model of transfer program complexity, and derives a number of results on program design, complexity and take up. Section 4 discusses the conceptual differences and similarities between productive complexity and "pure ordeals" (including a particular notion of stigma). Finally, section 5 offers some concluding remarks.

2 Classification Errors in Social Programs: The Existing Literature

We view non-participation by eligibles in social programs as a result of program parameters chosen by policy makers. Viewed in this way, it is natural to think of incomplete take up as a form of classification error of type I — a false negative. We introduce the following terminology:

Definition 1 (classification errors)

- *Type Ia errors (incomplete take up) occur if a program design results in some truly eligible individuals not applying for benefits.*
- *Type Ib errors (rejection errors) occur if a program design results in some truly eligible individuals applying for benefits and being rejected.*

[2] A case in point is the current discussion of health insurance reform in the United States, which stresses the importance of expanding health coverage as one of the main objectives of the reform.

- *Type II errors (award errors) occur if a program design results in some truly ineligible individuals applying for benefits and being accepted.*

For a government wanting to alleviate poverty among those who are truly eligible, being constrained by a limited budget, it is desirable to avoid all types of error. The occurrence of type Ia and type Ib errors undermine the goal of poverty alleviation, whereas the occurrence of type II errors make the program more expensive and divert government revenues away from other productive uses. Hence, the choice of parameters in a welfare program — benefits, eligibility rules and the complexity of the screening process — reflects the effect of each parameter on the different kinds of error. Indeed, a central message in this paper is that public programs have to be understood by integrating the treatment of all three types of classification error and considering the trade-off between them.

While a large empirical literature has analyzed incomplete take up and hence the occurrence of type Ia error, much fewer papers have estimated the occurrence of type Ib and type II errors. A small literature looking at classification error rates in U.S. Social Security disability award processes suggests that both award and rejection errors are very common. For example, the recent paper by Benítez-Silva et al. (2004) estimates the award error rate to about 20% and the rejection error rate to about 60%.[3] For the United Kingdom, Duclos (1995) studies classification errors in the Supplementary Benefits scheme, the main welfare program until 1988 providing means-tested cash benefits to the poor, and estimates an award error rate of 18.8% an a rejection error rate of 18.1%.

Opposite the empirical literature, a theoretical literature analyzes optimal transfer program design. Several strands of theoretical work are related to this paper. First, following on the seminal work of Mirrlees (1971), a large body of work studies the relationship between tax-transfer structures and the incentives for self-revelation. While much of this literature focuses on income taxation, a number of papers explicitly deals with the design of the social benefits (e.g. Diamond and Mirrlees, 1978; Nichols and Zeckhauser, 1982; Blackorby and Donaldson, 1988; Besley and Coate, 1992, 1995; Kanbur et al., 1994). Assuming that there exists no monitoring technology to assess true eligibility, this literature deals exclusively with type II errors and how to avoid them by restricting benefits in different ways. Some of this work focuses specifically on the role of ordeals in improving self-revelation. The notion of ordeals is related to the modeling of complexity in this paper, and we discuss the relationship between the two in Section 4.

Second, starting with the contribution by Akerlof (1978), a smaller strand of literature introduces a simplified monitoring technology into the mechanism design problem. This monitoring technology — labelled 'tagging' by Akerlof — can identify perfectly a given subset of eligibles. The screening process is imperfect because some eligibles are not tagged (a type Ib error) and it is exogenous to policy makers. While Akerlof did not allow for type II errors, subsequent work has incorporated two-sided classification error into the tagging framework, including papers by Stern (1982), Keen (1992), Diamond and Sheshinski (1995), Parsons (1996) and Salanié (2002). This literature assumes fixed award and rejection error

[3] An early study by Nagi (1969) reached broadly similar conclusions.

rates (exogenous monitoring technology), and it ignores incomplete take up. A recent paper by Jacquet and Van der Linden (2006) introduces stigma and imperfect take-up into the tagging model, keeping the assumption of an exogenous monitoring technology and ignoring the presence of type II errors.

Third, a much smaller literature considers monitoring in public sector programs, focusing on aspects such as agency problems among public sector workers (Boadway et al., 1999; Prendergast, 2007) and the monitoring of job search in unemployment insurance programs (Boadway and Cuff, 1999). The paper by Boadway et al. (1999) develops a tagging model in which the imperfectly observable effort of social workers affects the magnitude of type I and type II errors, and where costly monitoring is required to induce optimal effort among social workers. The paper characterizes the structure of the optimal tax-transfer system along with the optimal payment and monitoring of social workers.

Finally, overlapping with these different strands of literature, there is a large amount of work focusing on targeting transfers in the context of poverty-reduction programs in developing countries (see Coady et al., 2004, for a recent survey). In particular, this literature recognizes that eliminating Type II errors (better targeting of benefits) can adversely impact poverty reduction and cautions against use of indirect measures of transfer program targeting in assessing the performance of a program (see e.g., Ravallion, 2009, for a recent discussion).

Our paper contributes to the literature in two ways. First, we model and characterize the choice of complexity in social programs, accounting for the presence of imperfect take up in response to complexity. Second, we integrate the treatment of all three types of classification error, allowing for the magnitude of errors to be endogenous to program parameters chosen by policy makers. We show that optimal programs are typically characterized by all three types of error, consistent with the empirical evidence described above.

3 A Model of Social Program Complexity

3.1 Individuals

We assume that each individual is characterized by two parameters: an innate characteristic a and the precision by which this characteristic can be observed by outsiders σ. The characteristic a may reflect market productivity, or it may reflect other types of characteristics — say health or disability — depending on the program being considered. In the following, we refer to a simply as 'ability' or 'skill'. These skills are private information and cannot be ascertained directly by anyone else. Instead, if the individual attempts to claim welfare benefits, the government can test the individual and obtain a signal of true ability, $\tilde{a} = a + \varepsilon/\alpha$. The noise term ε reflects that program testing is imperfect, whereas the parameter α is a policy choice capturing the rigor of the test. We will come back to the interpretation and implications of α below.

Based on empirical analyses of benefit award processes (see Benítez-Silva et al., 2004), we assume that $\tilde{a}$ is a noisy but unbiased indicator of true ability so that ε is distributed with mean zero and variance σ^2. We assume that the normalized distribution of ε/σ (which has mean zero and variance one) is characterized by a c.d.f. $P(.)$, which is identical for everyone. We allow for the fact that the precision of measured skill, σ, may vary across individuals even if they have identical abilities. The heterogeneity in σ reflects that equally eligible individuals may test with more or less uncertainty in the welfare program. For example, aspects such as language barriers and unfamiliarity with the administrative procedures would create more uncertainty in the test and are heterogeneous across individuals. Moreover, the observability of true eligibility undoubtedly depends on occupation and opportunities to manipulate eligibility more generally, which creates additional heterogeneity in the noisiness of the eligibility indicator for individuals at the same earnings capacity. In other words, among otherwise equally eligible individuals (say, the low income population eligible for welfare benefits), there is heterogeneity in how they interact with the application process. We believe that this is a critical element for modeling imperfect take-up in a realistic way that matches the finding from the empirical literature that imperfect take-up cannot be easily explained by benefit structure and eligibility criteria.

As for the α-parameter, one possible interpretation is to view it as the number of tests. Under this interpretation, the government can subject an applicant to different tests in order to obtain indicators of skill. Each test leads to an indicator given by $a_i = a + \varepsilon_i$ where $\varepsilon_i \sim N(0, \sigma^2)$ — i.e., each indicator is a normally distributed unbiased indicator of the true skill level with variance σ^2. Examples of tests are interviews with case workers, a requirement to provide supporting documents, an opinion of a medical commission regarding disability, etc. Under this interpretation, the government estimates the skill of an individual using the arithmetic mean of α^2 tests, the distributional properties of which is exactly identical to $\tilde{a} = a + \varepsilon/\alpha$ with $\varepsilon \sim N(0, \sigma^2)$.[4]

More generally, the policy parameter α captures screening intensity and determines the extent of randomness in the application process. While increased screening intensity reduces randomness, it also creates complexity and imposes a burden on individuals. The cost of complexity can be interpreted as a cognitive cost of having a lot of testing, or it can be interpreted as a more concrete cost from the time and money spent applying. We represent the complexity cost of screening intensity α by a function $f(\alpha)$. It is important in the model that screening intensity α and complexity costs $f(\alpha)$ go hand in hand, so that program complexity is not fully unproductive (a so-called ordeal). In section 4, we discuss the desirability of unproductive complexity and its effect on the optimal amount of "productive" complexity.[5]

We assume that the government sets an eligibility criterion for receiving benefits denoted by $\bar{a}$. When the government relies on complexity α, benefits are granted to applicants who satisfy

$$\tilde{a} = a + \varepsilon/\alpha < \bar{a}. \tag{1}$$

The probability that an applicant with skill level a and precision σ receives benefits is

[4]Alternatively, if we had specified $\tilde{a} = a + \varepsilon/\sqrt{\alpha}$, then α (rather than α^2) would be the number of tests. We use the specification above because it is notationally simpler.

[5]The model implicitly allows for pure ordeal as a fixed cost component of $f(\alpha)$. One particular notion of stigma as representing exogenous psychic cost to an individual is consistent with this specification.

therefore given by $P\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right)$. We come back to the properties of $P(.)$ below.[6]

When making the participation decision, an individual knows the probability of being granted the benefit and trades off the potential utility gain from welfare payments against the cost of applying. We assume that utility depends on consumption C — equal to the sum of ability a and the (potential) welfare benefit B — and on application costs $f(\alpha)$. The utility level is given by $u(C - Af(\alpha))$, where A is an indicator variable for having applied. We make the standard assumption that $u(.)$ is increasing and weakly concave (allowing for the possibility of risk neutrality). We also assume that $\lim_{C\to\infty} u(C) = \infty$ and $\lim_{\alpha\to\infty} f(\alpha) = \infty$. An individual chooses to apply when

$$P\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right) u(a+B-f(\alpha)) + \left(1 - P\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right)\right) u(a-f(\alpha)) > u(a), \qquad (2)$$

and, conditional on applying, will receive benefits with the probability of $P\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right)$.

Ceteris paribus, a higher probability of receiving benefits increases the expected utility from applying. The probability of receiving benefits conditional on applying depends on the complexity parameter α, eligibility criterion $\bar{a}$, ability level a, and personal precision of ability signals σ. A higher $\bar{a}$ unambiguously increases the probability, whereas a higher a unambiguously decreases it. The effect of complexity α and precision σ depends on the sign of $\bar{a} - a$. When $\bar{a} > a$, higher complexity and better precision both increase the probability of receiving benefits. This is intuitive: when the individual is eligible under perfect information, reducing the noise in the eligibility metric is helpful. When $\bar{a} < a$, we have the opposite situation. While greater complexity may increase or decrease the likelihood of receiving benefits depending on the sign of $\bar{a} - a$, its effect on the ex post utility level is unambiguously negative regardless of whether benefits are received or not.

Using the participation constraint (2), we may solve for the minimum probability, $\tilde{P}$, consistent with applying for benefits:

$$\tilde{P}_a(\alpha, B) \equiv \frac{u(a) - u(a - f(\alpha))}{u(a+B-f(\alpha)) - u(a-f(\alpha))}. \qquad (3)$$

Individuals with a probability of receiving benefits above this critical value choose to apply for benefits, while the rest choose not to apply. In general, the threshold probability depends on the skill level a.[7] We will simplify the analysis by restricting attention to the class of preferences that eliminates the dependence of $\tilde{P}_a$ on a:

[6] Under the interpretation that the outcome of screening is based on a number of tests, this simple form of aggregating test results is potentially restrictive as it does not attempt to elicit information about the value of σ from the outcomes of the tests. However, while this is one possible interpretation of our model, we allow for more general interpretations. For example, it may be the case that, as the intensity of screening varies, the results still correspond to a single value of $\bar{a}$ and hence provide no additional information about σ. This could be the case if we interpret α as the amount of time an examiner devotes to the file of a program applicant.

[7] In particular, $\tilde{P}$ may be shown to be decreasing or increasing in ability depending on whether the utility function features decreasing or increasing absolute risk aversion. We do not have a strong prior as to whether higher ability individuals are willing to accept lower odds when applying for benefits, but the realistic case of decreasing absolute risk aversion would imply that this is the case. There are a number of other factors not modeled here that would have implications for this issue. For example, we restrict

Assumption 1 *The utility function has the Constant Absolute Risk Aversion (CARA) form,* $u(C) = \frac{1-e^{-\beta C}}{\beta}$*, where* $\beta \geq 0$ *(this specification reduces to risk-neutrality* $u(C) = C$ *for* $\beta = 0$*).*

Under this assumption, the threshold probability level for applying is given by

$$\tilde{P}(\alpha, B) = \begin{cases} \frac{1-e^{-\beta f(\alpha)}}{1-e^{-\beta B}}, \text{when } \beta > 0, \\ \frac{f(\alpha)}{B}, \quad \text{when } \beta = 0, \end{cases} \tag{4}$$

which is no longer a function of the ability level.[8] It is straightforward to show that $\frac{\partial \tilde{P}}{\partial \alpha} > 0$ and $\frac{\partial \tilde{P}}{\partial B} < 0$: a higher level of complexity increases the minimum acceptable probability of receiving benefits, whereas higher benefits decrease it.

Expressing the participation constraint as

$$P\left(\frac{\alpha(\bar{a} - a)}{\sigma}\right) > \tilde{P}(\alpha, B), \tag{5}$$

it can be solved for the precision level corresponding to indifference between applying and not applying:

$$\bar{\sigma}_a(\alpha, \bar{a}, B) \equiv \frac{\alpha(\bar{a} - a)}{P^{-1}\left(\tilde{P}(\alpha, B)\right)}. \tag{6}$$

When $\bar{a} > a$, individuals with σ lower than $\bar{\sigma}_a$ (high precision) apply for benefits. When $\bar{a} < a$, only individuals with σ greater than $\bar{\sigma}_a$ (low precision) choose to apply.

3.2 Population

We assume that there are two levels of ability: a low level a_L and a high level a_H. At each ability level, individuals are heterogeneous with respect to σ: for some, their ability level may be easily observable while for others it may be very difficult to ascertain without extensive testing. We note the following:

Remark 1 *At each ability level, let the precision of measured skill* σ *be distributed on* $[0, \infty)$*. There are three qualitative cases for the distribution of the probability of receiving benefits in the population:*

---

attention to a flat benefit although in practice the size of the benefit could depend on the realization of the indicator $\tilde{a}$. Letting the benefit depend negatively on $\tilde{a}$ would increase the minimum odds acceptable to the higher-ability individuals. Application costs may also vary with the ability level. On the one hand, if it is easier for high-ability applicants to file an application, their minimum acceptable probability would be lower. On the other hand, high-ability applicants tend to face higher opportunity cost of time spent applying, which would make their threshold probability higher.

[8] Although the CARA-assumption is not trivial, the only reason for making the assumption is that it eliminates the dependence of $\tilde{P}_a$ on a. We believe that this central implication of constant absolute risk aversion in the context of the model has an intuitive and realistic economic content: it implies that, given odds at which some low-ability individuals apply, we can always find a high-ability individuals who would also apply given the same odds. In our view, this is a realistic description of the real world, although it may reflect dimensions of heterogeneity not incorporated in our model. Hence, we would argue that a generalization of the model to non-CARA preferences should preserve this central property by adding additional dimensions of heterogeneity.

1. $\bar{a} \leq a_L < a_H$. Probabilities are in $(0, P(0)]$ and increasing in σ (always strictly increasing for high-ability individuals, strictly increasing for low-ability individuals only if $\bar{a} < a_L$).

2. $a_L < \bar{a} \leq a_H$. Probabilities for low-ability individuals are in $(P(0), 1)$ and strictly decreasing in σ; probabilities for high-ability individuals are in $(0, P(0)]$ and increasing in σ (strictly increasing if $\bar{a} < a_H$).

3. $a_L < a_H < \bar{a}$. Probabilities are in $(P(0), 1)$ for both types and are strictly decreasing in σ.

Whenever $a \neq \bar{a}$, any probability in the appropriate open interval, $(0, P(0))$ or $(P(0), 1)$, can be attained for some $\sigma \in [0, \infty)$.

This remark implies a "non-monotonicity" in committing Type II errors: they have to be committed when either $\bar{a} < a_L$ or $\bar{a} > a_H$, but not for intermediate values of $\bar{a}$. In the former cases, because the intervals of probabilities of receiving benefits are identical for the low- and high-ability populations, it will be impossible to avoid type II errors altogether. Note that $P(0)$ reflects a property of the normalized distribution of ε and therefore it is a constant independent of policy parameters or individual characteristics. In the natural case where the likelihoods of over- and understating true ability are identical such that median(ε) is zero, we have $P(0) = 1/2$. Recall that the threshold probability for applying, $\tilde{P}$, depends on the policy parameters α and B and, if these parameters are not constrained, $\tilde{P}$ can take any value. As a result,

Remark 2 *There exist policy parameters that result in no Type II errors ("full separation"); only low ability individuals apply. Such policies are characterized by $a_L < \bar{a} \leq a_H$ and $\tilde{P}(\alpha, B) \geq P(0)$.*

Moreover, there also exist policy parameters that additionally result in no Type Ia errors. They are characterized by $a_L < \bar{a} \leq a_H$ and $\tilde{P}(\alpha, B) = P(0)$.

One of the objectives of our analysis will be to determine whether policies with no type Ia and type II errors are optimal and, despite their apparent attractiveness, we will show that in the most interesting cases they are not. In particular, note that a government implementing a policy of full separation where $\tilde{P}(\alpha, B) = P(0)$ — i.e., no type II or type Ia errors — will continue to make Type Ib errors. That is, despite that only low-ability individuals are applying, some of them will be rejected. In fact, in the case of a symmetric distribution for ε where $P(0) = 1/2$, some low-ability applicants will face probabilities of receiving benefits as low as 1/2. Reducing the number of Type Ib errors can be accomplished by increasing the rigor of screening α, but in order to avoid Type Ia errors, the government must increase benefits correspondingly. Such increases are costly and, at the same time, constrained in their size when one wants to simultaneously discourage high-ability individuals from applying. As a result, the government faces serious constraints in pursuing policies that reduce the number of type Ib errors without introducing other types of classification error.

As we will demonstrate, these constraints may be severe enough to justify committing all three types of error.

To complete the characterization of the assumptions about the population, we need to specify the distribution of σ. We will denote the c.d.f. of the distribution of σ for ability-type a by G_a and the corresponding density function by g_a. The support of both distributions is assumed to be $[0, \infty)$. We assume that $g_a(0) = 0$, the density of individuals with perfectly observable skill is zero. The number of individuals of type a is given by $\bar{N}_a \equiv \int_0^\infty dG_a(\sigma)$, with both $\bar{N}_L$ and $\bar{N}_H$ assumed to be positive and finite.

Some of our results will depend on the following regularity assumptions:

Assumption 2 (thin tail for low ability) $\lim_{\sigma \to \infty} \sigma^2 g_L(\sigma) = 0$.

Assumption 3 (finite slope of density at zero for high ability) $\lim_{\sigma \to 0} g'_H(\sigma) < \infty$.

The first assumption states that the distribution of σ has no thick tail. In particular, it rules out the Pareto distribution, but it allows for distributions that have thinner tails such as the log-normal distribution. Intuitively, it will allow for the number of low-ability applicants to respond smoothly to policy changes that just discourage applying by everyone. The second assumption will guarantee that small changes in policy that make it beneficial for the high ability individuals to apply will result in only a small influx of them.

To summarize the model so far, Figure 1 illustrates the distribution of $P(.)$ and classification errors for a particular program based on a numerical simulation.[9] Although the specific shapes are not general, the figure illustrates the basic logic of the model. Both panels show the density of the P-distribution for low- and high-ability individuals, with Panel A highlighting the results for the low-types and Panel B highlighing the results for the high-types. The figure illustrates a program with $\bar{a} > a_H$, implying that $P(.)$ is distributed on the interval $(P(0), 1)$ for both types and is strictly decreasing in σ. We focus on this type of program, because it turns out to be interesting later on. The distribution of $P(.)$ is determined by the distribution of the noise term ε/σ (which is assumed to be normal so that $P(0) = 1/2$) along with the distribution of the precision of measured skill σ (which is assumed to be log-normal). Given $\bar{a} > a_H$, densities of $P(.)$ are positive everywhere in the open interval $\left(\frac{1}{2}, 1\right)$ for both types, because any probability in this interval can be attained for some $\sigma \in [0, \infty)$. At a given σ, low-ability applicants have a higher probability of being awarded benefits, and hence the P-distribution for low-ability individuals is shifted to the right compared to the distribution for high-ability individuals. The two types have the same threshold probability $\tilde{P}$ $(= 0.736$ in the simulation), and individuals with higher Ps than this (corresponding to those with low σs) apply for benefits. The program is associated with all three types of classification error. In the low-ability distribution, type Ia (take-up) errors are committed in the region to the left of $\tilde{P}$, while type Ib errors occur in the region to the right because probabilities of acceptance are lower than 1. In the high-ability distribution, type II errors occur in the region to the right of $\tilde{P}$ because probabilities of acceptance are

[9] See the note under the figure for the parameter values. The working paper version (Kleven and Kopczuk, 2008) contains extensive simulations.

greater than zero (in fact, greater than 0.736). An interesting question is whether a program outcome of this kind can be an equilibrium outcome. To study this question, we turn to the final piece of the model: the specification of the government's objective.

3.3 Government

We consider a problem of income maintenance extending the specification of Besley and Coate (1992, 1995). They considered the design of income maintenance programs ensuring that each individual obtains a target minimum benefit at a minimum fiscal cost. In our model (as in reality), we do not necessarily have full participation, because eligible individuals may choose not to apply for the benefit and because eligible applicants may be rejected by program administrators due to imperfect testing. Hence, the objective becomes to provide a minimum benefit for as many low-ability (truly deserving) individuals as possible, being constrained by a limited budget.[10]

Denoting by $\bar{B}$ the target minimum benefit and by R the exogenously given budget size, the government's problem may be written as

$$\max_{\alpha,\bar{a},B} \quad N_L\left(\alpha,\bar{a},B\right) \tag{7}$$

subject to

$$\left[N_L\left(\alpha,\bar{a},B\right)+N_H\left(\alpha,\bar{a},B\right)\right]B \leq R \tag{8}$$

and

$$B \geq \bar{B}, \tag{9}$$

where

$$N_a(\bar{a},\alpha,B) = \int_0^{\bar{\sigma}_a(\alpha,\bar{a},B)} P\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right) d\,G_a(\sigma) \tag{10}$$

is the number of successful applicants of type a as a function of policy parameters.

There are several aspects of our policy objective that deserve mentioning. First, the income maintenance goal implies that policy makers are not directly concerned with the utility cost that program complexity imposes on individuals (i.e., apart from its effect on benefit take-up). As mentioned in the beginning, we consider income maintenance instead of a social welfare maximization in order to simplify the proofs and results. In addition, the income maintenance approach arguably fits better with actual political debates than social welfare maximization. If politicians care about utility instead of income, it may seem obvious that complexity is a less effective instrument. However, it is important to note that complexity in our model is not just an ordeal (a pure deadweight cost on applicants), but a byproduct of efforts to screen between deserving and undeserving applicants. Indeed, because the ordeal-part of complexity imposes a higher utility cost on low-ability applicants than on high-ability applicants (due to concave utility), the screening benefits of complexity

[10]Interestingly, as noted by Besley and Coate (1992), this policy objective fits Mill's (1848) characterization of the poverty-alleviation problem as "how to give the greatest amount of needful help, with the smallest encouragement to undue reliance on it."

are driven entirely by the second effect in our model.[11] This sets our paper apart from the large literature on ordeals (e.g., Nichols and Zeckhauser, 1982; Besley and Coate, 1992, 1995; Cuff, 2000; Kreiner and Tranæs, 2005). While it is often difficult to justify pure ordeals without resorting to non-Paretian objectives (see Kreiner and Tranæs, 2005, for a discussion), complexity in our model continues to be associated with desirable screening effects in the context of social welfare maximization. In Appendix B, we outline a welfarist approach to complexity and discuss how this extension would affect our findings.

Second, we assume that benefits cannot fall below some minimum value despite that, in general, not all of the low-ability individuals are going to receive benefits (note though that the government can increase benefits above $\bar{B}$). Reducing benefits to a small enough value would allow for providing benefits to everyone, and therefore allowing for unrestricted benefits is incompatible with a non-trivial problem of maximizing the number of deserving recipients. Absent a direct welfarist objective, providing a target minimum income to successful recipients is a natural way of modeling the goal of poverty alleviation.

Third, we do not model the revenue side of the system. While the distortions introduced are undoubtedly important, our model does not necessarily describe the full society. Rather, our "high-ability" individuals should be viewed as still relatively poor but not poor enough to be in need of social welfare. Under this interpretation, benefits are financed by a wealthier (and not modeled) segment of the society. While the exogeneity of budget size R with respect to program design and outcomes may be questioned, there is a wide class of models in which, at any given budget size, policy parameters should be selected optimally conditional on that budget. Hence, an alternative way of interpreting our approach is as a solution to a subproblem of this kind.

Fourth, our model ignores the administrative costs of various policy instruments. It is natural to expect, for example, that higher values of α are costlier to implement and that the administrative cost of the program may vary with the number of applicants and recipients. While these are undeniably important factors in practice (which should be analyzed in follow-up work), ignoring them makes the model more tractable and allows us to focus on the trade-offs between different instruments that result from their effects on the different types of classification error.

Fifth, the government pursues policies that are horizontally inequitable as some low-ability individuals are going to receive benefits while others will not. It is not possible to pursue a horizontally equitable policy unless one is able to provide benefits to everyone — rich and poor. This is a property of this model and, likely, of the real world: in order to reach every poor individual we would have to accept a very large number of Type II errors.[12]

In the following section, we characterize social programs that solve the problem specified above. We show that the solution depends, among other things, on the size of the program budget R. We restrict attention to program budgets satisfying $R < \bar{B}\left(\bar{N}_L + \bar{N}_H\right)$. If the

[11] In Section 4 we argue that pure ordeals are not desirable in our model.

[12] Policies that are horizontally inequitable also arise naturally in the work on imperfect tagging (for example, Akerlof, 1978; Parsons, 1996).

budget were larger than this, the government's problem has a simple solution: the number of low-ability applicants reaches its theoretical maximum $\bar{N}_L$ by giving a universal benefit $B \geq \bar{B}$ to everybody, which is an affordable policy when $R \geq \bar{B}\left(\bar{N}_L + \bar{N}_H\right)$. A universal benefit would be implemented by letting the eligibility criterion $\bar{a}$ tend to infinity, in which case the probability of receiving benefits tends to 1 for everybody.

3.4 Results

We begin our analysis of the model by specifying the first-best allocation that the government would pursue under full information.

Definition 2 (first best) *Suppose that it is possible to observe both a and σ. Then the optimal policy provides benefits of at least $\bar{B}$ to $\min\left(R/\bar{B}, \bar{N}_L\right)$ individuals with ability a_L (and $B = \bar{B}$ when $R/\bar{B} \leq N_L$). The choice of these individuals is undetermined (there may be many first-best policies).*

The requirement that a first-best program must reach $\min\left(R/\bar{B}, \bar{N}_L\right)$ recipients amounts to saying that the program either spends the entire budget R or, if there are unused funds, this is because there are no low-ability individuals left who have not received benefits. Notice that this definition of a first-best policy is conditional on the *exogenous* funds R allocated to the program, and therefore does not account for the fact that the amount of revenue allocated to a program may in itself depend on the information available to policy makers. In particular, if information were perfect, it would not make sense to allocate funds to a program such that $R > \bar{B} \cdot \bar{N}_L$ given the problem specified in (7)-(9). The purpose of the above definition is not to specify a "global" first best, but to specify the best possible outcome at any given budget size R in order to create a benchmark against which to compare the actual outcome for a program designed under imperfect information. In the presence of imperfect information, even if a program has a large budget, $R > \bar{B} \cdot \bar{N}_L$, it will not be able to reach all low-ability individuals, and this is therefore a relevant case to consider. We come back to this point below.

As noted in Remark 2, there exist policies that result in providing benefits only to low-ability individuals. In certain cases, it is possible to achieve one of the first-best allocations despite the lack of perfect information.

Proposition 1 (first best) *First-best programs always involve full separation.*[13] *For R small enough, first-best is feasible and the optimal program is characterized by $B = \bar{B}$, $a_L < \bar{a} \leq a_H$, and $\tilde{P}(\alpha, \bar{B}) \geq P(0)$. The optimum is not necessarily unique.*

Proof. Setting policy instruments such that $a_L < \bar{a} \leq a_H$, $B = \bar{B}$, and $\tilde{P}(\alpha, \bar{B}) \geq P(0)$ ensures that (i) benefits are provided only to low-ability individuals and (ii) each recipient receives only the target minimum. The final requirement for a program to be first best is that $N_L = \min\left(R/\bar{B}, \bar{N}_L\right)$. To see that this is only possible if the budget is "small", notice that the class of programs specified above can never reach all low-ability individuals. The number of low-ability recipients within this class of programs is maximized by setting $\bar{a} = a_H$. Given $\bar{a} = a_H$ and $B = \bar{B}$, it is not possible to

[13] Here and everywhere else in the analysis, we of course limit attention to the case when universal programs are not feasible, $R < (\bar{N}_L + \bar{N}_H)\bar{B}$

set α such that $N_L = \bar{N}_L$. This is because, at any finite α, the probability of rejection for each low-ability applicant, $1 - P\left(\frac{\alpha(a_H - a_L)}{\sigma}\right)$, is greater than zero, and α cannot be increased without bound because the associated increase in $\tilde{P}(\alpha, \bar{B})$ would ultimately discourage all applications. Hence, there is a maximum number of low-ability individuals $N_L^* < \bar{N}_L$ that can be reached within the class of programs we have specified. Define $R^* = \bar{B} \cdot N_L^* < \bar{B} \cdot \bar{N}_L$ as the largest budget that can be spent on this type of program, and denote by α^* the level of α that achieves N_L^*. Now, for $R > R^*$, we always have $N_L < \min\left(R/\bar{B}, \bar{N}_L\right)$ and therefore not first best. Conversely, for any $R < R^*$, we can always ensure $N_L = \min\left(R/\bar{B}, \bar{N}_L\right)$ by increasing α beyond α^* (because a higher α increases $\tilde{P}$ and low-ability recipients fall to zero for high enough α because complexity costs outweigh benefits). Finally, because first-best programs are always associated with $R/\bar{B} < \bar{N}_L$, and because programs that provide benefits to any high-ability individual imply $N_L < R/\bar{B}$, first-best policies always involve full separation. ■

The proposition shows that, if the budget is small enough, we can spend all of the money providing the target minimum $\bar{B}$ to low-ability individuals only, which is the first-best outcome at the given budget. In particular, this is the case for $R \leq R^*$ where $R^* < \bar{B} \cdot \bar{N}_L$. At the other extreme, if the budget is very large, $R \geq \bar{B}\left(\bar{N}_L + \bar{N}_H\right) \equiv \bar{R}$, we pointed out above that the optimal program offers a universal benefit $B \geq \bar{B}$ to everybody. The most interesting case is the one in between the small-budget case $R \in (0, R^*]$ and the very-large-budget case $R \in \left(\bar{R}, \infty\right]$, i.e. where $R \in \left(R^*, \bar{R}\right)$, in which case first best is not feasible. The rest of this section is devoted to characterizing optimal social programs in this intermediate range.

We have to consider both full separation (but non-first best) programs and non-full separation programs. We start by noting that, at any budget size, full separation policies are always feasible:

Lemma 1 (type II errors can be avoided) *At any budget size, there exists a policy that satisfies the budget constraint and involves full separation.*

Proof. We just need to consider the situations where the first-best policy is not feasible. Let us consider policies involving $\tilde{P}(\alpha, B) = P(0)$ and $\bar{a} = a_H$. For such policies, no high-ability individuals apply, whereas all of the low-ability individuals apply. Consider increasing α while simultaneously increasing B to keep $\tilde{P}(\alpha, B) = P(0)$. By construction, this policy will retain full separation. As $\alpha \to \infty$, $P\left(\frac{\alpha(\bar{a} - a_L)}{\sigma}\right)$ increases and tends to one for any σ and therefore, because all of the low-ability individuals apply, the number of low-ability individuals receiving benefits increases and tends to $\bar{N}_L$. Simultaneously, $\alpha \to \infty$ and $\tilde{P}(\alpha, B) = P(0)$, implies that $B \to \infty$ and therefore spending will be tending to ∞. Hence, at some point the full budget will be spent. ■

Notice that, when the budget is large, full-separation policies that exhaust the entire budget feature benefits that are higher than $\bar{B}$. Higher benefits attract high-ability individuals, but they can be discouraged from applying by having a high degree of complexity.

To characterize the optimal policy under full separation, we will need the following lemma:

Lemma 2 *Consider $a < \bar{a}$ and $\tilde{P}(\alpha, B) = P(0)$. Under assumption 2,*

1. *a small increase in α increases the number of individuals receiving benefits $N_a(\alpha, \bar{a}, B)$ (even though it reduces the number of applicants)*

2. *a small decrease in B has no effect on the number of individuals receiving benefits $N_a(\alpha, \bar{a}, B)$*

3. *$N_a(\alpha, \bar{a}, B)$ is continuously differentiable in α and B (despite switching from everyone applying to non-full take-up).*

Proof. See the appendix. ■

Assumption 2 guarantees smoothness of the number of beneficiaries when α and B change so as to just stop some people from applying: these are the people with the highest variances and the thin-tail assumption implies that there are not 'many' of them. We can then characterize the optimal full-separation, non-first best program as follows:

Proposition 2 (best policy avoiding type II errors) *Under assumption 2, the best policy implementing full separation when the first-best allocation is not feasible is characterized by $B > \bar{B}$, $\bar{a} = a_H$, $\tilde{P}(\alpha, B) > P(0)$, and $\frac{\partial N_L}{\partial \alpha} = 0$.*

Proof. First, since the first-best allocation is not feasible, a full separation policy that spends all of the budget must have $B > \bar{B}$. By Lemma 1, there exist full separation policies that satisfy the budget constraint.

Second, the best full separation policy involves $\bar{a} = a_H$. To see this, suppose instead that $\bar{a} < a_H$ in the optimum. Then we can increase $\bar{a}$ to a_H, which would imply more low-ability people receiving benefits. Now, we are spending too much money, but we can reduce B until the budget is satisfied (this is possible because initially $B > \bar{B}$ and at $\bar{B}$ not everything is spent). In the new equilibrium, we have $N_L B = R$ and a lower B, so that N_L must be higher, contradicting that $\bar{a} < a_H$ was optimal.

Third, the optimal policy involves $\tilde{P}(\alpha, B) > P(0)$. Conversely, suppose that $\tilde{P}(\alpha, B) = P(0)$. Consider increasing α slightly so that $\tilde{P}(\alpha, B) > P(0)$. By Lemma 2, the number of low-ability recipients increases and spending increases over R. Therefore, we may now reduce B until spending falls to R (we may do so because $B > \bar{B}$ to begin with). We end up with all of the budget spent, lower benefits and therefore more low-ability individuals receiving benefits — a contradiction.

Finally, having established that $B > \bar{B}$, $a = a_H$, and $\tilde{P}(\alpha, B) > P(0)$, the problem is to maximize $N_L(\alpha, a_H, B)$ with respect to α and B, subject to $N_L(\alpha, a_H, B)B = R$. The latter equation can be solved for $B = B(\alpha)$ where $\frac{\partial B}{\partial \alpha} = -\frac{B\frac{\partial N_L}{\partial \alpha}}{N_L + B\frac{\partial N_L}{\partial B}}$. The problem we solve is now equivalent to maximizing $N_L(\alpha, a_H, B(\alpha))$ with respect to α. The first-order condition is $\frac{\partial N_L}{\partial \alpha} + \frac{\partial N_L}{\partial B}\frac{\partial B}{\partial \alpha} = 0$. Substituting for $\frac{\partial B}{\partial \alpha}$ and rearranging and simplifying yields $\frac{\partial N_L}{\partial \alpha} N_L = 0$ so that $\frac{\partial N_L}{\partial \alpha} = 0$ (note that we have to have $N_L > 0$ to satisfy $BN_L = R$). We are guaranteed that $\frac{\partial N_L}{\partial \alpha} = 0$ has a solution, because N_L is positive and increasing in α at $\tilde{P}(\alpha, B) = P(0)$ (by Lemma 2), N_L is equal to zero when $\tilde{P}(\alpha, B) = 1$, and $\tilde{P}(\alpha, B)$ itself increases with α (and attains the value of one for a sufficiently high α). ■

This proposition has several implications. First, because $\tilde{P}(\alpha, B) > P(0)$ both kinds of Type I error are made:

Corollary 1 *The best policy that avoids Type II errors involves both Type Ia and Ib errors.*

Although the objective is to maximize the number of low-ability recipients and the government is able to discourage high-ability individuals from applying, the best full-separation policy is associated with incomplete take up. The reason is that discouraging high-ability individuals from applying makes it impossible to provide benefits to all of the low-ability individuals who do apply. Given that some Type Ib errors are being made, it is always optimal to reduce their number somewhat at the cost of introducing some Type Ia errors.

Second, the optimal full-separation policy features benefits that are higher than the minimum required level $\bar{B}$. This is a mechanical result. Given that full separation imposes a restriction on $\bar{a}$ and given that a sufficiently high α discourages applications, the only way to spend all of the budget while retaining full separation is by increasing B.

Third, the optimal full separation policy involves setting complexity α such that it has no effect at the margin on the number of low-ability recipients. This implies that the additional discouragement of low-ability applicants from a higher complexity cost (operating through $\tilde{P}(\alpha, B)$) is exactly offset by a higher probability of receiving benefits conditional on applying, $P\left(\frac{\alpha(\bar{a}-a_L)}{\sigma}\right)$. As we shall see below, this result does not carry over to optimal non-separation policies.

Finally, observe that full-separation policies (whether first-best is feasible or not) may be very costly in terms of the complexity burden that they impose on welfare recipients. As an example, consider the case of risk-neutrality where $\tilde{P}(\alpha, B) = \frac{f(\alpha)}{B}$. At the optimum, we have $\tilde{P}(\alpha, B) \geq P(0)$ and therefore $f(\alpha) \geq P(0)B$. Hence, the cost of complexity consumes at least a fraction $P(0)$ of welfare transfers. Recall that $P(0)$ is the probability that an individual will test below his true ability level. Under the natural assumption that the distribution of tests is symmetric, i.e. $P(0) = \frac{1}{2}$, complexity consumes at least one-half of the income surplus for those who get the benefit. Since some applicants are rejected in the process, aggregate complexity costs may then constitute more than half of the surplus to all welfare applicants.

So far, we have imposed the rigid restriction that the policy maker attempts to keep high-ability individuals from applying. This must be the best policy if one can simultaneously set $B = \bar{B}$ because the number of the low-ability recipients then reaches its theoretical maximum. However, as we have shown, this is possible only if the budget is small enough (Proposition 1). For greater budgets, the best full separation policy requires overpaying benefits (Proposition 2), which suggests that it may be optimal to let some high-ability individuals into the program, while simultaneously reducing benefits to allow for more recipients in total. Indeed, we can show

Lemma 3 *Under assumption 3, we can improve upon the policy characterized in Proposition 2 by increasing $\bar{a}$ slightly.*

Proof. In the appendix. ■

This result follows because a small increase in the eligibility threshold above a_H has only a second-order effect on the number of high-ability recipients who are just becoming eligible, while having a first-order effect on the number of low-ability recipients. This allows for reducing the benefit below the level prevailing under the optimal full-separation policy (where $B > \bar{B}$), and therefore financing a higher number of low-ability recipients. Hence,

Corollary 2 (type II errors are optimal) *When the first-best allocation cannot be implemented, the second-best policy always involves non-separation.*

This is an important result. Even though it is possible to discourage high-ability individuals from applying, it is not optimal. The optimal policy will therefore involve both Type I and Type II errors.

The rest of this section will be devoted to characterizing the optimal policy under non-full separation. While Lemma 3 establishes that there exists non-separation programs with $\bar{a} > a_H$ that dominate the best full separation program, we have to consider the possibility that the *optimal* non-separation program is associated with a "stringent" eligibility criterion, i.e. $\bar{a} \leq a_H$. All else equal, a stringent eligibility criterion will of course discourage high-ability applicants from applying, but we can bring them back in by having a low complexity so that $\tilde{P}(\alpha, B) < P(0)$. However, we can show that non-separation policies combining a stringent eligibility criterion with low complexity ($\bar{a} \leq a_H$ and $\tilde{P}(\alpha, B) < P(0)$) are always dominated by non-separation policies that combine a lenient eligibility criterion with high complexity ($\bar{a} > a_H$ and $\tilde{P}(\alpha, B) \geq P(0)$):

Proposition 3 (eligibility criterion is "lenient") *When the first-best allocation is not feasible, setting $\bar{a} \leq a_H$ is never optimal.*

Proof. Lemma 3 implies that, if the first best is not feasible, the full separation policy is not optimal. Therefore, we want to consider non-full separation policies where $\bar{a} \leq a_H$ and $\tilde{P} < P(0)$. Suppose a policy of this kind, denoted by $(\alpha^*, \bar{a}^*, B^*)$, is optimal. Consider then an alternative policy that keeps $B = B^*$, sets $\bar{a}$ to satisfy $\max\{\bar{a}^*, a_L\}$, and increases α to obtain $\tilde{P} = P(0)$. The number of high-ability applicants drops to zero, whereas all of the low-ability applicants will apply with the probability of receiving benefits increasing for each of them.[14] This change therefore increases the value of the objective function. If this policy results in a reduction in the total number of beneficiaries (note that all of the previous high-ability recipients drop out), it is affordable and therefore it is an improvement — contradiction. Otherwise, if the policy increases the total number of recipients, it is not affordable. If $B^* > \bar{B}$, we can then reduce benefits. Such an adjustment will maintain full separation and if it yields an affordable policy it must be an improvement because full budget will be spent on lower benefits paid to low-ability individuals only. This again contradicts the optimality of the original policy. When reducing benefits to $\bar{B}$ results in a policy that is still unaffordable, that implies that it is possible to spend more than the full budget on a first-best allocation and therefore the first-best allocation can be implemented as in the proof of Proposition 1, thereby contradicting the assumption that first-best allocation is not feasible. ■

The intuition for the result in Proposition 3 is that stringent programs with low complexity are associated with a lot of type Ib errors. In this case, a combined increase in the eligibility criterion $\bar{a}$ and screening intensity α is very effective, because it allows full separation between the two types along with higher probabilities of receiving benefits for all low-ability applicants. If this policy change maintains affordability, we have a welfare improvement; if it does not maintain affordability, we would be in the first-best case where all the funds can be spent providing the target minimum benefit to low-ability individuals only.

The following lemma demonstrates that the optimal policy changes smoothly from the first-best region to the non-full separation region.

Lemma 4 *Denote by R^* the maximum budget that allows for implementing the first-best allocation and let $(\alpha, \bar{a}, B) = \left(\alpha^*, a^H, \bar{B}\right)$ be the corresponding optimal policy. Denote by $x(R) = (\alpha(R), \bar{a}(R), B(R))$ the optimal policies as a function of R, $R \geq R^*$. The function $x(R)$ is right-continuous at R^*.*

[14] For the case where $\max\{\bar{a}^*, a_L\} = a_L$, there is a technical qualification due to the fact that the participation constraint (5) is written with strict inequality. Because of this, if $\bar{a} = a_L$ and $\tilde{P} = P(0)$, the low-ability individuals would not apply. To be precise, the government would instead have to set $\bar{a} = a_L + \delta$ where δ can be arbitrarily small, in which case all low-ability individuals would apply.

Proof. In the appendix. ■

Recall the structure of our problem: we maximize the number of low-ability recipients N_L subject to the constraints $(N_L + N_H)B = R$ and $B \geq \bar{B}$. Lemma 2 guarantees that N_L and N_H are both continuously differentiable at $\tilde{P}(\alpha, B) = P(0)$ which is the only point where it is not immediately obvious. Therefore, the maximum satisfies the following first-order conditions (where λ is the Lagrange multiplier associated with the government budget):

$$\frac{\partial N_L}{\partial \alpha} - \lambda \left[\frac{\partial N_L}{\partial \alpha} + \frac{\partial N_H}{\partial \alpha}\right] B = 0, \tag{11}$$

$$\frac{\partial N_L}{\partial \bar{a}} - \lambda \left[\frac{\partial N_L}{\partial \bar{a}} + \frac{\partial N_H}{\partial \bar{a}}\right] B = 0, \tag{12}$$

$$\left[\frac{\partial N_L}{\partial B} - \lambda \left(\frac{\partial N_L}{\partial B} + \frac{\partial N_H}{\partial B}\right) B - \lambda \left(N_L + N_H\right)\right] (B - \bar{B}) = 0, \tag{13}$$

where the first bracketed term in equation (13) is non-positive and the second is non-negative. When one considers a restricted problem of selecting α and $\bar{a}$ holding B constant, condition (13) need not hold but equations (11) and (12) remain valid. Consequently, as long as α and $\bar{a}$ are selected optimally given B, we must have:

$$\frac{\partial N_L / \partial \bar{a}}{\partial N_H / \partial \bar{a}} = \frac{\partial N_L / \partial \alpha}{\partial N_H / \partial \alpha} = \frac{\lambda B}{1 - \lambda B}\,. \tag{14}$$

Neither eligibility criterion $\bar{a}$ nor the intensity of screening/complexity α have a direct revenue cost. Therefore, intuitively, what matters in comparing them is how well each of them screens low- from high-ability individuals. This is summarized by the marginal change in the number of low-ability recipients relative to the marginal change in the number of high-ability recipients. At the optimum, the two instruments screen equally well. It is also straightforward to show that when $\frac{\partial N_L/\partial \bar{a}}{\partial N_H/\ \partial \bar{a}} > \frac{\partial N_L/\partial \alpha}{\partial N_H \partial \alpha}$, $\bar{a}$ should be increased and/or α reduced, with the opposite implication when the sign of this inequality is reversed.

In general, the effect of α on the number of recipients of each type, $\frac{\partial N_L}{\partial \alpha}$ and $\frac{\partial N_H}{\partial \alpha}$, may be either positive or negative. But because $\frac{\partial N_L/\partial \bar{a}}{\partial N_H/\partial \bar{a}}$ is always positive, any program satisfying eq. (14) must be associated with complexity such that $\frac{\partial N_L}{\partial \alpha}$ and $\frac{\partial N_H}{\partial \alpha}$ have the same sign.[15]

Before continuing, we state the following very useful identity that links derivatives of the number of recipients with respect to the three instruments (the proof is in the appendix):

$$\frac{\partial N_a}{\partial \alpha} = \frac{\bar{a} - a}{\alpha} \frac{\partial N_a}{\partial \bar{a}} + \frac{\partial \tilde{P}/\partial \alpha}{\partial \tilde{P}/\partial B} \frac{\partial N_a}{\partial B}, \quad \text{where} \quad \frac{\partial \tilde{P}/\partial \alpha}{\partial \tilde{P}/\partial B} = -\frac{e^{\beta B} - 1}{e^{\beta f(\alpha)} - 1} \cdot f'(\alpha) \tag{15}$$

This result follows from the fact that all three instruments operate through two margins. First, instruments can affect $\tilde{P}(\alpha, B)$, the minimum acceptable probability of receiving benefits consistent with applying. Second, instruments can affect the actual probability of receiving benefits for those who apply, $P\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right)$. Complexity works through both

[15] In numerical simulations that we performed, the optimal solution was associated with $\frac{\partial N_L}{\partial \alpha}$ and $\frac{\partial N_H}{\partial \alpha}$ being negative.

margins, whereas benefits work only through the first one and the eligibility criterion works only through the second one.

We can now show that the government pursues social policies associated with incomplete take up:

Proposition 4 (Type Ia errors are optimal) *Under assumption 2, when the first-best allocation is not feasible, for any value of B, the optimal choice of $\bar{a}$ and α implies $\tilde{P}(\alpha, B) > P(0)$.*

Proof. Conversely, suppose that $\tilde{P}(\alpha, B) \leq P(0)$. Then, we have $\frac{\partial N_L}{\partial B} = 0$ and $\frac{\partial N_H}{\partial B} = 0$ (Lemma 2 shows that this holds at $\tilde{P}(\alpha, B) = P(0)$, while it obviously holds at $\tilde{P}(\alpha, B) < P(0)$). As a consequence, identity (15) becomes $\frac{\partial N_a}{\partial \alpha} = \frac{\bar{a}-a}{\alpha} \frac{\partial N_a}{\partial \bar{a}}$, which implies

$$\frac{\partial N_L / \partial \alpha}{\partial N_H / \partial \alpha} = \frac{\bar{a} - a_L}{\bar{a} - a_H} \frac{\partial N_L / \partial \bar{a}}{\partial N_H / \partial \bar{a}} > \frac{\partial N_L / \partial \bar{a}}{\partial N_H / \partial \bar{a}}$$

That, however, implies that the original policy could not have been optimal because it violates the optimality condition (14) (and, in fact, α should be increased). ∎

Corollary 2 and Proposition 4 together imply that, when the budget is not small, social programs feature both type Ia and type II errors. Optimal programs of course also feature type Ib errors because, given that α and $\bar{a} > a_H$ are finite, low-ability applicants face probabilities of receiving benefits $P\left(\frac{\alpha(\bar{a}-a_L)}{\sigma}\right)$ distributed on $\left(\tilde{P}, 1\right)$, and hence face non-zero probabilities of rejection.[16] We have therefore shown that large-budget programs are associated with all three types of classification error as illustrated in Figure 1 discussed earlier.

An interesting issue regarding the optimal setting of instruments is the choice of the optimal level of benefits. Given that the government cares only about the number of recipients, it may seem obvious that benefits should be set at the lowest possible level. But recall that the best full-separation policy did not have this property: according to Proposition 2, benefits should be increased above their minimum level. In that context, this was a mechanical result driven by the inability to otherwise spend all of the budget on eligibles. However, notice that benefits also play a screening role by potentially attracting high- and low-ability applicants at different rates. Intuitively, if benefits are sufficiently good at screening, this may warrant increasing them despite their budgetary cost. This possibility is reinforced by the next proposition.

Proposition 5 (B screens better than α and $\bar{a}$) *Suppose that α and $\bar{a}$ are set optimally given B. Then,*

$$\frac{\partial N_L / \partial B}{\partial N_H / \partial B} > \frac{\partial N_L / \partial \bar{a}}{\partial N_H / \partial \bar{a}} = \frac{\partial N_L / \partial \alpha}{\partial N_H / \partial \alpha} = \frac{\lambda B}{1 - \lambda B}. \tag{16}$$

[16] Notice that α cannot increase without bound because in that case, to prevent everybody from dropping out of the program due to prohibitive application costs, B would also have to increase without limit, which requires an unlimited budget, $R = \infty$. The eligibility criterion $\bar{a}$ also cannot increase without bound unless the budget is large enough to give benefits to everyone, i.e. $R \geq \left[\bar{N}_L + \bar{N}_H\right] \bar{B}$, in which case a universal program would be optimal.

Proof. Equalities in the statement of the proposition repeat equation (14). To show that the inequality is valid, recall identity (15) and note that $\tilde{P}(\alpha, B)$ does not depend on the type. Therefore,

$$\frac{\frac{\partial N_L}{\partial B}}{\frac{\partial N_H}{\partial B}} = \frac{\frac{\bar{a}-a_L}{\alpha}\frac{\partial N_L}{\partial \bar{a}} - \frac{\partial N_L}{\partial \alpha}}{\frac{\bar{a}-a_H}{\alpha}\frac{\partial N_H}{\partial \bar{a}} - \frac{\partial N_H}{\partial \alpha}} = \frac{\frac{a_H-a_L}{\alpha}\frac{\partial N_L}{\partial \bar{a}}}{\frac{\bar{a}-a_H}{\alpha}\frac{\partial N_H}{\partial \bar{a}} - \frac{\partial N_H}{\partial \alpha}} + \frac{\frac{\bar{a}-a_H}{\alpha}\frac{\partial N_L}{\partial \bar{a}} - \frac{\partial N_L}{\partial \alpha}}{\frac{\bar{a}-a_H}{\alpha}\frac{\partial N_H}{\partial \bar{a}} - \frac{\partial N_H}{\partial \alpha}}$$

The denominator of the first term is equal to $-\frac{\partial \tilde{P}/\partial \alpha}{\partial \tilde{P}/\partial B}$ times $\frac{\partial N_H}{\partial B}$ and it is positive because $\partial \tilde{P}/\partial \alpha > 0$ while $\partial \tilde{P}/\partial B < 0$. Therefore, the first term of the expression is unambiguously positive. When α and $\bar{a}$ are set optimally given B, we have $\frac{\partial N_L/\partial \bar{a}}{\partial N_H/\partial \bar{a}} = \frac{\partial N_L/\partial \alpha}{\partial N_H/\partial \alpha} = \frac{\lambda B}{1-\lambda B}$ and it is straightforward to show that in this case the second term is equal to $\frac{\lambda B}{1-\lambda B}$. Therefore, we have

$$\frac{\frac{\partial N_L}{\partial B}}{\frac{\partial N_H}{\partial B}} = \frac{\frac{a_H-a_L}{\alpha}\frac{\partial N_L}{\partial \bar{a}}}{-\frac{\partial \tilde{P}/\partial \alpha}{\partial \tilde{P}/\partial B}\frac{\partial N_H}{\partial B}} + \frac{\lambda B}{1-\lambda B} > \frac{\lambda B}{1-\lambda B}. \tag{17}$$

■

This is an important result: on the margin, benefits are better at screening low- from high-ability individuals than the other two instruments (complexity and eligibility). This is a global result that holds for *any* value of B as long as the other instruments are set optimally — it does not assume that B itself has been optimized and it holds regardless of how B is financed. Therefore, the only reason not to increase benefits beyond $\bar{B}$ is the revenue cost. In other words, the question is whether the advantage from using benefits to screen is big enough to compensate for the extra revenue cost. From the screening point of view, increasing benefits is preferred to the other two instruments.

What is the intuition for the strong screening power of benefits? Technically, the result follows from identity (15) that links the screening effects of benefits, eligibility criteria and complexity. To understand the logic here, observe that the number of recipients is determined by two factors: the minimum odds necessary to apply and the probability of receiving benefits conditional on applying, as embedded in the participation constraint (5). Benefits affect the first margin, eligibility criteria affect the second one, while complexity affects both. The effect of complexity may then be decomposed into benefit (or rather cost) and eligibility effects. Absent its cost, increasing complexity (tightening screening) works better than modifying eligibility criteria: this is reflected in condition (15) by the factor $\bar{a} - a$ multiplying the effect of eligibility on the number of recipients. This factor is greater for low ability applicants than for high ability applicants. At the optimum then, if both eligibility and complexity are equally good at screening, the screening implications of the cost part of complexity must be biased against low-ability applicants. Equivalently though, it implies that benefits have to be particularly good at distinguishing the two types when the other instruments are optimized.

While the details of this discussion are model-specific, the main observation here is that the effect of complexity combines the intensity of screening effect and the cost effect, with the intensity of screening working better than eligibility in distinguishing the two types and the cost effect working in the same way as benefits do (though, obviously, with the opposite sign). These two observations imply that benefits have to be particularly good at screening. Both of them follow from our assumptions but are likely to be true more generally: it is intuitive to expect that increasing intensity of screening is better than modifying eligibility criteria because it magnifies visible differences between types and, while the effect of benefits and costs of applying need not be symmetric, there is no obvious presumption in which direction the asymmetry, if any, should go.

When $\bar{a}$ and α are set optimally, benefits should be increased from some level B if

$$\frac{\partial N_L}{\partial B} - \lambda B\left(\frac{\partial N_L}{\partial B} + \frac{\partial N_H}{\partial B} + \frac{N_L + N_H}{B}\right) > 0 \quad \Rightarrow \quad \frac{\frac{\partial N_L}{\partial B}}{\frac{\partial N_H}{\partial B}} > \frac{\lambda B}{1-\lambda B} + \frac{\lambda B}{1-\lambda B}\frac{N_L + N_H}{B\frac{\partial N_H}{\partial B}}$$

Recall equation (14): $\frac{\lambda B}{1-\lambda B}$ reflects the optimal extent of screening performed by the other instruments. Benefits should be used beyond their minimal level only if they are sufficiently better than the other instruments at screening by a factor identified in the last term — this is a correction for the budgetary cost of increasing benefits. It is difficult for benefits to satisfy this condition if the other instruments are already good at screening ($\frac{\lambda B}{1-\lambda B}$ is high), when there are a lot of individuals whose benefits will have to be increased ($N_L + N_H$ is high) and when $B\frac{\partial N_H}{\partial B}$ is small. Substituting for $\frac{\partial N_L/\partial B}{\partial N_H/\partial B}$ using the equality in (17) yields

$$\frac{a_H - a_L}{-\alpha\frac{\partial \tilde{P}/\partial \alpha}{\partial \tilde{P}/\partial B}}\frac{\partial N_L}{\partial \bar{a}} > \frac{\lambda B}{1-\lambda B}\frac{N_L + N_H}{B}.$$

Recalling that $\frac{\partial N_L/\partial \bar{a}}{\partial N_H/\partial \bar{a}} = \frac{\lambda B}{1-\lambda B}$, we can rewrite this to

$$\frac{a_H - a_L}{\alpha}\frac{\partial N_H}{\partial \bar{a}} > -\frac{\partial \tilde{P}/\partial \alpha}{\partial \tilde{P}/\partial B}\frac{N_L + N_H}{B} = -\frac{\partial \tilde{P}/\partial \alpha}{\partial \tilde{P}/\partial B}\frac{R}{B^2}, \tag{18}$$

where the last equality uses the budget identity $R = B(N_L + N_H)$. It is optimal to increase B beyond $\bar{B}$ if this condition holds when evaluated at $\bar{B}$ and the optimal α and $\bar{a}$ (at $\bar{B}$). How should we interpret this condition? B should be used if changes in eligibility $\bar{a}$ bring too many high-ability individuals, where the inequality gives the specific meaning to "too many". Alternatively, note from identity (15) that $\frac{\bar{a}-a_L}{\alpha}\frac{\partial N_H}{\partial \bar{a}}$ is the effect of α on the number of high-ability recipients while holding the number of applicants constant (i.e. holding $\tilde{P}$ constant). Thus, the same condition can be expressed in terms of extra complexity bringing in "too many" high-ability individuals. The following proposition characterizes the optimal choice of benefits.

Proposition 6 (optimal benefits) *Suppose that the first-best allocation is not feasible, $R \geq R^*$. Denote by N_L^* the number of low-ability recipients under the best full separation policy identified by Proposition 2. Let $\bar{a}^* = \inf_{\bar{a}}\left\{\max_{\alpha} N_L(\alpha, \bar{a}, \bar{B}) \geq N_L^*\right\}$, $a_H < \bar{a}^* < \infty$.*[17] *Then,*

1. *For R sufficiently close to R^*, setting $B = \bar{B}$ is optimal.*
2. *For R sufficiently large, setting $B = \bar{B}$ is optimal.*

[17]We know that $\bar{a}^* > a_H$ because, from the proof of Proposition 3, if the first-best allocation is not feasible, any non-full separation policy with $\bar{a} \leq a_H$ is dominated by a full separation policy and therefore also dominated by the best full separation policy N_L^*.

3. *A sufficient condition for $B > \bar{B}$ is given by*

$$\frac{\bar{a}^* - a_H}{\bar{a}^* - a_L} G_L^{-1}(N_L^*) \geq G_H^{-1}\left(\frac{1}{P(0)}\left(\frac{R}{\bar{B}} - N_L^*\right)\right). \tag{19}$$

Proof. In the appendix. ■

The intuition for the first result is straightforward. When the budget is small (but large enough to make first best infeasible), the eligibility threshold $\bar{a}$ will be very close to a_H. As we cross a_H with $\bar{a}$, initially we are still mostly providing benefits to low-ability individuals (on the margin, the share of high-ability recipients is close to zero when $\bar{a}$ is close to a_H). Given the presence of such a good instrument that does not have a direct revenue cost, it must dominate any instrument that does have a revenue cost (such as B). That is, as the number of high-ability applicants is initially small, any screening benefits of using high benefits have to be dominated by the costly nature of this instrument. The second part is also intuitive: as the budget size increases, the number of individuals served increases as well and therefore increasing benefits becomes more costly.

Part 3 is the most interesting. It gives the sufficient condition for $B > \bar{B}$. Moreover, note that this condition can be satisfied by varying G_H without affecting any of the other variables in eq. (19). In particular, the definitions of N_L^* and $\bar{a}^*$ are based solely on the low-ability distribution and parameters — they do not depend on G_H. Furthermore, the argument of G_H^{-1} on the right-hand side depends on constants R, $\bar{B}$ and again on N_L^* so that it is independent of G_H. Finally, note that both the left-hand side and the argument of G_H^{-1} are positive due to the fact that $\bar{a}^* > a_H$ and N_L^* delivers fewer low-ability recipients than $R/\bar{B}$, which is what the first-best policy would deliver. Thus, given parameters and low-ability distribution, we will be able to find some distributions G_H that satisfy the condition identified in the above proposition. All that is required is selecting the distribution so that there are more than $\frac{1}{P(0)}(R/\bar{B} - N_L^*)$ high-ability individuals with σ smaller than $\frac{\bar{a}^* - a_H}{\bar{a}^* - a_L} G_L^{-1}(N_L^*)$. This is a requirement imposed on G_H at a particular strictly positive point.[18] The corollary below is a consequence of this reasoning and it highlights that the case $B > \bar{B}$ cannot be dismissed as being irrelevant because it will apply when the number of high-ability applicants is sufficiently large.

Corollary 3 *Fix the parameters of the problem other than the high-ability distribution. Select some distribution of high-ability individuals $G_H^0(\sigma)$ (with the corresponding number of high-ability individuals N_H^0) and consider a class of distributions $G_H^\eta(\sigma) = \eta G_H^0(\sigma)$ (with the corresponding number of high-ability individuals ηN_H^0). For high enough η, setting $B > \bar{B}$ is optimal.*

Proof. For sufficiently high η,

$$G_H^\eta\left(\frac{\bar{a}^* - a_H}{\bar{a}^* - a_L} G_L^{-1}(N_L^*)\right) = \eta G_H^0\left(\frac{\bar{a}^* - a_H}{\bar{a}^* - a_L} G_L^{-1}(N_L^*)\right) \geq \frac{1}{P(0)}\left(\frac{R}{\bar{B}} - N_L^*\right).$$

and this condition is equivalent to the inequality in Part 3 of Proposition 6. ■

[18] In particular, there is no restriction imposed on the properties of $G_H(.)$ around zero so that we can pick a distribution satisfying Assumption 3.

Intuitively, as the number of high-ability individuals grows, the importance of screening increases so that government should rely on benefits that have been shown to be particularly good at screening despite their adverse revenue implications.

4 Productive Complexity versus Pure Ordeals

In this section, we briefly discuss two different notions of complexity costs that the model allows for. Our focus so far has been on the cost of complexity varying as a function of the screening parameter α, but one might also consider changing the level of complexity costs conditional on the value of α. The effect of changes in α on the cost may be thought of as "productive complexity" — this is a by-product of an attempt to screen better. Changes in the cost given α may be interpreted as pure ordeals — they are not related to the extent of screening but reduce the welfare of a program applicant, holding other things constant. In our context, it is natural to think of such ordeal costs as "unproductive complexity", but they may include other types of costs such as fixed psychic costs of program participation due to stigma. As an illustration, we consider a simple formulation of the cost of complexity where $f(\alpha) = s + c \cdot \alpha$. In this example, s may be interpreted as the pure ordeal while $c \cdot \alpha$ is the cost of productive complexity.

If complexity were a pure ordeal, it would be able to serve a screening function only if it imposed a higher utility cost on high-ability individuals than on low-ability individuals, conditional on income (e.g. Nichols and Zeckhauser, 1982). This condition is not satisfied in our model where the utility cost of complexity is always higher for low-ability individuals under the assumed preference specification. This suggests that introducing pure ordeals or unproductive complexity in social programs is not socially optimal in the context of the model.[19] The role of pure ordeals can be more precisely understood by considering the participation constraint (5): a change in s affects individual behavior only through the minimum odds required for applying, $\tilde{P}(\alpha, B)$. Given the independence of $\tilde{P}$ from a and σ, this effect is the same for everyone and equivalent to reducing the value of B. We have shown in the previous section that benefits are particularly good at screening (to such an extent that they may even be increased above the minimum even despite their revenue cost) and hence the same argument implies that stigma is particularly bad at screening. As a result, if s were a policy parameter, it should be set to zero.[20]

To shed further light on the effect of various types of complexity costs, a longer version of the paper (Kleven and Kopczuk, 2008) considers numerical simulations that vary the pure ordeal and complexity parameters s and c.[21] In the simulations, an increase in pure ordeals leads to a more lenient eligibility criterion in order to compensate for the fact that individuals are less willing to apply. Pure ordeals have two offsetting effects on complexity:

[19]These comments on the use of ordeals are specific to the specification of preferences in our model. It is an empirical question whether stigma and complexity impose lower or higher costs on high-ability individuals than on low-ability individuals.

[20]Of course, another difference between s and B in the model is that the latter has revenue cost. If increasing s was costly, the conclusion would naturally be unchanged.

[21]Here we view the ordeal-parameter s as a primitive of the model (say, a fixed psychic cost of stigma) rather than a policy parameter that is optimized.

(1) Higher ordeals discourage all individuals from program participation, which calls for a compensating reduction in complexity-driven transaction costs. (2) Those who continue to apply have higher probabilities of receiving benefits on average and contain a larger share of undeserving individuals. It then becomes harder to discriminate between low- and high-ability individuals, which makes it necessary to increase the rigor of testing. In our simulations, the second effect weakly dominates at low levels of the fixed cost, while the first effect strongly dominates at higher levels resulting in a non-monotonic relationship between the complexity parameter α and pure ordeals s.

While ordeals have a strong negative impact on equilibrium take up among the intended recipients, they have an indirect desirable effect on the amount of rejection errors. The reduction in rejection errors may seem beneficial by itself, but it is a consequence of the difficulty to screen in the presence of ordeals: at very high levels of ordeals, virtually all applicants of both types are approved. Moreover, higher ordeals increase the number of high-ability applicants and create more type II errors in equilibrium, because of the slackening of policy instruments $\bar{a}$ and α along with the fact that pure ordeals are less costly for the non-deserving individuals. Overall, ordeals are bad for targeting efficiency as they reduce the number of deserving recipients and increase the number of undeserving recipients.[22]

As a result, the model predicts that programs with high pure ordeals (or high stigma) should be characterized by low take up, almost no rejection errors, and relatively low productive screening. This prediction arguably fits the empirical pattern: welfare programs such as TANF and Food Stamps tend to be associated with high stigma, low take up, low rates of false rejection and arguably a relatively low degree of screening, whereas social insurance programs like DI are characterized by low stigma, high take up, high rates of false rejection and a very high degree of complexity.

5 Concluding Remarks

This paper stresses the importance of complexity in public programs for the take up of benefits. The empirical literature on benefit take up shows that transaction costs arising from complex application processes constitute important barriers to enrollment in social programs. A large part of the observed complexity is related to screening efforts such as eligibility criteria, documentation, forms, interviewing, testing, and appeals, which suggests that complexity is not a pure deadweight cost (an ordeal) imposed on applicants. We therefore model complexity as a by-product of efforts to screen between deserving and undeserving applicants, and study the choice of screening intensity and associated complexity by policy makers interested in poverty alleviation. In the characterization of program design, we integrate the study of take up with the study of classification errors of type I (rejection errors) and type II (award errors). While a more rigorous screening technology has desirable effects on the amount of rejection and award errors, the associated complexity introduces transaction costs into the application process and induces incomplete take up.

We find that optimal programs that are not universal always feature a high screening

[22]This finding is consistent with the earlier argument that, if ordeals were a policy instrument, it should not be used for screening purposes in the context of our framework.

intensity and complexity in order to restrict the number of rejection and award errors. Although complexity hurts take up among the truly deserving, the government can use eligibility and benefit instruments to partially compensate for this. Hence, the model can explain the presence of high complexity and incomplete take-up as equilibrium outcomes of policy making under imperfect information. It is central to this result that complexity is driven by screening efforts and is therefore not fully unproductive. Our model also incorporates unproductive complexity (an ordeal), and we argued in Section 4 that this form of complexity is bad for targeting efficiency and should not be used as a policy instrument under the assumptions of our model.

The model may be able to shed light on the heterogeneity of complexity, targeting, take up and classification errors across different social programs. We find that program design and outcomes depend on the following key parameters: the maximum number of individuals that can be reached (i.e., budget size relative to the target benefit), the number of potential applicants, the distribution of true skill, the distribution of measured skill, complexity costs, and pure ordeal costs. The fact that all of these parameters are likely to vary simultaneously across different social programs complicates a comparison of the theoretical predictions and real-world programs. A careful calibration exercise could facilitate such an analysis, but is beyond the scope of this paper.

Our model makes a number of simplifying assumptions that we would like to discuss briefly. First, we consider policy makers who are interested in income maintenance and therefore not directly concerned with the utility cost that complexity imposes on program applicants. We outline an extension of the model to social welfare maximization in Appendix B. While this extension introduces some additional and interesting effects, we argue that the model retains the central insights of the simpler income maintenance model. In particular, the basic screening argument for complexity studied in the context of the income maintenance model carry over to the welfarist model. The welfarist approach introduces two additional offsetting effects on complexity and take up. On the one hand, by placing a weight on the utility cost of complexity, the approach calls for less complexity and higher take up. On the other hand, rejection errors are worse than take-up errors under the welfarist approach, because the former is associated with a utility cost of applying. This favors policies that can substitute take-up errors for rejection errors, which creates an argument for more complexity and lower take up.

Second, while the model accounts for complexity costs borne by individuals, it ignores that complexity may also create program costs borne by the government. Accounting for administrative costs will have offsetting effects on equilibrium complexity. On the one hand, allowing for complexity to be positively associated with administrative costs obviously makes it less effective as a policy instrument. On the other hand, it seems reasonable that administrative costs would also depend positively on the number of applications that have to be processed, which favors instruments that are capable of increasing the number of deserving recipients with no accompanying increase in the number of applicants. This tends to improve the efficacy of complexity as it reduces the amount of rejection errors while making it more costly for would-be applicants to claim the benefit.

The model makes several other simplifying assumptions such as focusing on a two-type ability distribution and restricting attention to a flat benefit that does not depend on the eligibility signal. Extending the model in these directions would enrich the analysis and possibly add empirical content to the theory. This paper should be seen as an initial step towards incorporating administrative complexity as a choice variable in policy analysis. Although several authors have suggested that complexity is an important aspect of policy design, for example in the context of tax policy (Slemrod, 1990; Slemrod and Yitzhaki, 2002), little has been done in terms of actual modeling. We have emphasized an application to the design of transfer programs and the take up of social benefits, but our model may be applicable to other areas of public policy.

A Proofs

Proof of Lemma 2. We want to evaluate $\frac{\partial N_a}{\partial \alpha}$ and $\frac{\partial N_a}{\partial B}$ when $\bar{a} > a$ and $\tilde{P}(\alpha, B) = P(0)$. Begin with the first of these. Differentiating formula (10) explicitly yields

$$\frac{\partial N_a}{\partial \alpha} = \int_0^{\bar{\sigma}_a(\alpha,\bar{a},B)} p(\cdot)\frac{\bar{a}-a}{\sigma}dG_a(\sigma) + \frac{\partial \bar{\sigma}_a}{\partial \alpha}\tilde{P}(\alpha, B)g_a(\bar{\sigma}_a). \tag{A.1}$$

The first term is unambiguously positive. We will show that the second term is zero when $\bar{a} > a$ and $\tilde{P}(\alpha, B) = P(0)$. Evaluating $\frac{\partial \bar{\sigma}_a}{\partial \alpha}$ yields

$$\frac{\partial \bar{\sigma}_a}{\partial \alpha} = \frac{(\bar{a}-a)P^{-1} - \alpha(\bar{a}-a)\left(p\left(\frac{\alpha(\bar{a}-a)}{\bar{\sigma}_a}\right)\right)^{-1}\frac{\partial \tilde{P}}{\partial \alpha}}{\left(P^{-1}\left(\tilde{P}(\alpha,B)\right)\right)^2} = \frac{\bar{\sigma}_a}{\alpha} - \frac{\bar{\sigma}_a^2}{\alpha(\bar{a}-a)}\left(p\left(\frac{\alpha(\bar{a}-a)}{\bar{\sigma}_a}\right)\right)^{-1}\frac{\partial \tilde{P}}{\partial \alpha}. \tag{A.2}$$

Note that $\lim_{\tilde{P}(\alpha,B)\to P(0)} \bar{\sigma}_a = \infty$: as the threshold probability of receiving benefits approaches $P(0)$, the number of individuals not applying goes to zero. All other terms in expression (A.2): α, $\bar{a}-a$, $\frac{\partial \tilde{P}}{\partial \alpha}$ and $p(\cdot)$ are positive and finite (the argument of $p(\cdot)$ goes to 0 as $\bar{\sigma}_a$ goes to infinity and density at 0 is positive). Consequently, as we approach $P(0)$ with $\tilde{P}$, expression (A.2) tends to $-\infty$ at the rate of $\bar{\sigma}^2$. Consequently, the behavior of the second term of (A.1) depends on the behavior of $\bar{\sigma}_a^2 g(\bar{\sigma}_a)$ and, by assumption 2, $\lim_{\bar{\sigma}_a\to\infty} \bar{\sigma}_a^2 g(\bar{\sigma}_a) = 0$.

Now consider $\frac{\partial N_a}{\partial B}$. It is equal to $\frac{\partial \bar{\sigma}_a}{\partial B}\tilde{P}(\alpha, B)g_a(\bar{\sigma}_a)$ and we have $\frac{\partial \bar{\sigma}_a}{\partial B} = -\frac{\alpha(\bar{a}-a)}{(P^{-1}(\tilde{P}))^2}\frac{dP^{-1}(\tilde{P})}{dB} = -\bar{\sigma}_a^2\frac{1}{\alpha(\bar{a}-a)}\frac{dP^{-1}(\tilde{P})}{dB}$. It is straightforward to show as before that all terms but $\bar{\sigma}_a$ and $g(\bar{\sigma}_a)$ are bounded away from zero and infinity. Therefore $\frac{\partial N_a}{\partial B}$ behaves as $g(\bar{\sigma}_a)\bar{\sigma}_a^2$ and thus it is zero in the limit by assumption 2.

Finally, the second term of (A.1) is uniformly zero when $\tilde{P}(\alpha, B) < P(0)$ and the first term is continuous. Therefore, the whole expression in (A.1) is continuous, which proves the third part of the lemma. Similarly, $\frac{\partial N_a}{\partial B}$ is uniformly zero when $\tilde{P}(\alpha, B) < P(0)$, so that it is also continuous at $\tilde{P}(\alpha, B) = P(0)$. ■

Proof of Lemma 3. Denote by (α^*, a_H, B^*) the best policy under full-separation characterized in Proposition 2, and consider increasing $\bar{a}$ above a_H. We will show first that the right-derivative of N_H with respect to $\bar{a}$ is equal to zero at (α^*, a_H, B^*): $\frac{\partial N_H(\alpha^*, a_H, B^*)}{\partial \bar{a}_+} = 0$. Differentiating (10) with respect to $\bar{a}$ yields

$$\frac{\partial N_a}{\partial \bar{a}} = \int_0^{\bar{\sigma}_a(\alpha,\bar{a},B)} p\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right)\frac{\alpha}{\sigma}g_a(\sigma)d\sigma + \frac{\partial \bar{\sigma}_a}{\partial \bar{a}}\tilde{P}(\alpha, B)g_a(\bar{\sigma}_a), \tag{A.3}$$

such that

$$\frac{\partial N_H(\alpha^*, a_H, B^*)}{\partial \bar{a}_+} = \lim_{\substack{\bar{a}\to a_H \\ \bar{a}>a_H}}\left\{\int_0^{\bar{\sigma}_H(\alpha^*,\bar{a},B^*)} p(\cdot)\frac{\alpha^*}{\sigma}g_H(\sigma)d\sigma + \frac{\partial \bar{\sigma}_H}{\partial \bar{a}_+}\tilde{P}(\alpha^*, B^*)g_H\left(\bar{\sigma}_H(\alpha^*, \bar{a}, B^*)\right)\right\}.$$

Note that $\lim_{\bar{a}\to a_H} \bar{\sigma}_H = 0$, and therefore the second term is zero: we have $g(0) = 0$ while the other components $\tilde{P}$ and $\frac{\partial \bar{\sigma}_H}{\partial \bar{a}_+} = \frac{\alpha}{P^{-1}(\cdot)}$ tend to finite limits (the limit of $P^{-1}(\cdot)$ is positive because we are considering a point with $\tilde{P} > P(0)$). In the first term, the integrand $p(\cdot)\frac{\alpha^*}{\sigma}g_H(\sigma)$ is bounded from above in the neighborhood of $\sigma = 0$ by assumption 3 and because $p(\cdot)$ is bounded from above. Therefore, the first term tends to zero.

On the other hand, for the low-ability individuals we have that $a_L < \bar{a}$, $\bar{\sigma}_L > 0$, and $\frac{\partial \bar{\sigma}_L}{\partial \bar{a}} > 0$, so that the derivative $\frac{\partial N_L}{\partial \bar{a}}$ given by A.3 is strictly positive.

Hence, starting at the best full-separation policy (α^*, a_H, B^*), we can increase the threshold $\bar{a}$ slightly above a_H so as to give benefits to more low-ability people, while bringing in only an infinitesimal number of high-ability people. We would then be spending too much money, but we can reduce B below B^* until the revenue constraint is satisfied. At this new equilibrium, since B is lower, the total number of recipients, $N_L + N_H$, is higher. Moreover, since the number of high-ability recipients is infinitesimal, the number of low-ability recipients is higher than before. ■

Proof of Lemma 4. Consider $R = R^* + \varepsilon$. Denote the optimal policy at R^* by $(\alpha^*, a^H, \bar{B})$ and note that it involves $\tilde{P} > P(0)$ (if $\tilde{P} = P(0)$, increasing the number of low-ability recipients could be increased by increasing α by Lemma 2 and therefore implement the first-best for an even greater budget). To simplify notation, denote the optimal policy given ε by $(\alpha(\varepsilon), \bar{a}(\varepsilon), B(\varepsilon))$.

We will show first that $\lim_{\varepsilon\to 0} \bar{\sigma}_H = 0$. To see that note that the proof of Proposition 2 implies that we can achieve $\frac{R^*+\varepsilon}{B} > N_L > \frac{R^*}{B}$ by sticking to the full separation policies. Note that we must have $N_H < \frac{\varepsilon}{B}$ because otherwise $N_L \geq \frac{R^*}{B}$ and the optimal non-separation policy would be preferred. By definition $N_H = \int_0^{\bar{\sigma}_H} P\left(\frac{\alpha(\bar{a}-a_H)}{\sigma_H}\right) dG_H(\sigma_H)$. We know that $P\left(\frac{\alpha(\bar{a}-a_H)}{\sigma_H}\right) > P(0)$ for everyone because $\bar{a} > a_H$. Therefore, $N_H > G_H(\bar{\sigma}_H)P(0)$ and consequently $G_H(\bar{\sigma}_H) < \frac{N_H}{P(0)} < \frac{\varepsilon}{\bar{B}P(0)}$, implying that $\lim_{\varepsilon\to 0} \bar{\sigma}_H = 0$.

Now observe that $\lim_{\varepsilon\to 0} \bar{\sigma}_H = 0$ implies $\lim_{\varepsilon\to 0} \bar{a} = a_H$. Recall that $\bar{\sigma}_H = \frac{\alpha(\bar{a}-a_H)}{P^{-1}(\tilde{P}(\alpha,B))}$ and $\bar{\sigma}_L = \frac{\alpha(\bar{a}-a_L)}{P^{-1}(\tilde{P}(\alpha,B))}$ so that $\bar{\sigma}_H = \frac{\bar{a}-a_H}{\bar{a}-a_L}\bar{\sigma}_L$. Therefore, $\bar{a} - a_H = \frac{\bar{\sigma}_H}{\bar{\sigma}_L - \bar{\sigma}_H}(a_H - a_L)$. Note also that $G_L(\bar{\sigma}_L) > \frac{R^*}{B}$ (the number of low-ability applicants which is still greater than the number of recipients must be at least as high as in the full-separation optimum). Consequently, as $\bar{\sigma}_H \to 0$, $\frac{\bar{\sigma}^H}{\bar{\sigma}^L - \bar{\sigma}^H} \to 0$ and therefore $\bar{a} - a_H \to 0$.

Consider what happens when $\varepsilon \to 0$. The resulting value of the objective is $N_L(\varepsilon)$. N_L is a continuous function of policy parameters in the relevant region.[23] Suppose that $\lim_{\varepsilon\to 0}(\alpha, \bar{a}, B) \neq (\alpha^*, \bar{a}^*, B^*)$. In that case, $\lim_{\varepsilon\to 0} N_L(\varepsilon) = N_L(\lim_{\varepsilon\to 0} \alpha(\varepsilon), \lim_{\varepsilon\to 0} \bar{a}(\varepsilon), \lim_{\varepsilon\to 0} B(\varepsilon)) < N_L(\alpha^*, \bar{a}^*, B^*)$ — the last inequality follows from the fact that the limiting point $(\lim_{\varepsilon\to 0} \alpha(\varepsilon), \lim_{\varepsilon\to 0} \bar{a}(\varepsilon), \lim_{\varepsilon\to 0} B(\varepsilon))$ implements full separation because $\tilde{P}(\alpha(\varepsilon), B(\varepsilon)) > P(0)$ (by Proposition 4), $a(\varepsilon) \to a_H$ as demonstrated earlier and $(\alpha^*, \bar{a}^*, B^*)$ was the optimal point under full separation. This is however a contradiction because it implies that for sufficiently small ε we would have been better off using the full separation policy (and not using all of the money). Consequently, $\lim_{\varepsilon\to 0}(\alpha, \bar{a}, B) = (\alpha^*, a^H, \bar{B})$. ■

Proof of identity 15. Recall the definition of N_a, equation (10) and the definition of $\bar{\sigma}_a$ in equation (6). α affects N_a through two channels. First, $\alpha(\bar{a} - a)$ is the maximum realization of the individual error term that results in receiving benefits — it enters both the integrand in N_a and the limit of integration. Second, the minimum acceptable probability threshold $\tilde{P}(\alpha, B)$ influences who applies. The effect of α on N_a is the sum of these two effects. Instrument $\bar{a}$ works only on the first margin, while instrument B works only on the second margin. Recognizing that allows for writing the effect of α as a combination of the effects with respect to the other two probabilities. ■

Proof of Proposition 6. Part 1.

Let α^* be the optimal value of α at the full separation policy with maximum budget R^*. Denote by $\tilde{\lambda}$ the Lagrange multiplier from the problem of maximizing the objective function with respect to α and $\bar{a}$ while setting $B = \bar{B}$. It can be easily shown that we will want to increase B over $\bar{B}$ if and only if

$$\frac{\partial N_L}{\partial B} > \frac{\tilde{\lambda}B}{1 - \tilde{\lambda}B}\frac{\partial N_H}{\partial B} + \frac{\tilde{\lambda}B}{1 - \tilde{\lambda}B}\frac{N_L + N_H}{B} \tag{A.4}$$

The left-hand side is non-negative and we don't have to worry about it increasing without bounds as ε changes — it converges to a finite limit of $\frac{\partial N_L}{\partial B}(\alpha^*, a_H, \bar{B})$. All terms on the right-hand side are non-negative and $\frac{N_L+N_H}{B}$ is finite and bounded away from zero ($N_L > \frac{R^*}{B}$ and $B = \bar{B}$). We will show that $\frac{\tilde{\lambda}B}{1-\tilde{\lambda}B} \to \infty$ as $\varepsilon \to 0$ and thus this inequality is violated for small enough ε. To see that, recall that $\frac{\tilde{\lambda}B}{1-\tilde{\lambda}B} = \frac{\partial N_L/\partial\bar{a}}{\partial N_H/\partial\bar{a}}$. We will show that the numerator is finite while the denominator falls

[23] N_a is continuous when $\bar{a} > a$ and has a discontinuity at $\bar{a} = a$ when $\tilde{P} = P(0)$. In this case, the discontinuity is at a_L, but we are considering $\bar{a} \geq a_H > a_L$.

to zero. To see that, write explicitly $\partial N_a/\partial\bar{a}$:

$$\frac{\partial N_a}{\partial \bar{a}} = \int_0^{\bar{\sigma}_a} \alpha p\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right)\frac{g_a(\sigma)}{\sigma}d\sigma + \frac{\alpha}{P^{-1}(\tilde{P}(\alpha,B))}\tilde{P}(\alpha,B)g_a(\bar{\sigma}_a) \tag{A.5}$$

All terms here are non-negative. The first-term vanishes for the high-types as $\varepsilon \to 0$, because $\lim_{\varepsilon\to 0}\bar{\sigma}_H = 0$ while the integrand is bounded away from infinity in the neighborhood of $\sigma = 0$ by Assumption 3. It remains positive for low-ability individuals because $\bar{\sigma}_L$ remains bounded away from zero as $\varepsilon \to 0$. By Lemma 4, $\lim_{\varepsilon\to 0}\tilde{P}(\alpha,B) = \tilde{P}(\alpha^*,\bar{B}) > P(0)$, so that $P^{-1}(\tilde{P}(\alpha,B))$ has non-zero limit. Consequently, for the high-ability types the second term disappears because $g_H(0) = 0$ while it remains positive for the low-ability types. As a result $\lim_{\varepsilon\to 0}\frac{\tilde{\lambda}B}{1-\tilde{\lambda}B} = \lim_{\varepsilon\to 0}\frac{\partial N_L/\partial\bar{a}}{\partial N_H/\partial\bar{a}} = \infty$, implying that the inequality (A.4) must be violated for sufficiently small ε and therefore $B = \bar{B}$ is optimal for sufficiently small εs.

Part 2.

This is an implication of condition (18) evaluated at $\bar{B}$ and the optimal α and $\bar{a}$. To see this, hold $\bar{B}$ constant and increase R. The parameter α is bounded by $P(0) < \tilde{P}\left(\alpha,\bar{B}\right) < 1$, so that $\frac{\partial\tilde{P}/\partial\alpha}{\partial\tilde{P}/\partial B} = -\frac{e^{\beta B}-1}{e^{\beta f(\alpha)}-1}f'(\alpha)$ is bounded away from zero when evaluated at the optimal α and $\bar{B}$.

Moreover, we can show that $\frac{\partial N_H}{\partial\bar{a}} \to 0$ as we keep increasing R. To see this, start by noting that, as R increases, $\bar{a}$ has to increase. There exists a finite budget size $\bar{R} \geq \bar{B}\left(N_L^* + N_H^*\right)$ at which everyone receives benefits, and at that budget size we have $\bar{a} = \infty$ and $\bar{\sigma}_a = \infty$. As R approaches this value, we have $\bar{a} \to \infty$ and $\bar{\sigma}_a \to \infty$.

Now recall eq. (A.5) and consider what happens as $\bar{a}$ and $\bar{\sigma}_a$ increases. The second term can be written as $\tilde{P}(\alpha,B)\frac{\bar{\sigma}_a g_a(\bar{\sigma}_a)}{(\bar{a}-a)}$. We must have $\lim_{\bar{\sigma}_a\to\infty}\bar{\sigma}_a g_a(\bar{\sigma}_a) = 0$ (if the limit exists), because otherwise $g_a(\cdot)$ would not be a distribution function. Moreover, $\bar{a}-a$ tends to ∞ so that the second term disappears in the limit. For the first term, integration by parts yields

$$\int_0^{\bar{\sigma}_a} \alpha p\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right)\frac{g_a(\sigma)}{\sigma}d\sigma = \frac{1}{\bar{a}-a}\left\{-P\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right)\sigma g_a(\sigma)\Big|_0^{\bar{\sigma}_a} + \int_0^{\bar{\sigma}_a} P\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right)\left[g_a(\sigma)+\sigma g_a'(\sigma)\right]d\sigma\right\}$$

$$= \frac{1}{\bar{a}-a}\left\{\left(1-\tilde{P}(\alpha,B)\right)\bar{\sigma}_a g_a(\bar{\sigma}_a) + \int_0^{\bar{\sigma}_a} P\left(\frac{\alpha(\bar{a}-a)}{\sigma}\right)\left[g_a(\sigma)+\sigma g_a'(\sigma)\right]d\sigma\right\}.$$

We have that $\frac{1}{\bar{a}-a}$ tends to zero. The first-term in the bracket disappears as $\bar{\sigma}_a$ tends to infinity. Because $P(\cdot)$ is a c.d.f., it can be bounded from above by 1, so that the second term is smaller than $\int_0^{\bar{\sigma}_a} g_a(\sigma) + \sigma g_a'(\sigma)d\sigma = \sigma g_a(\sigma)\Big|_0^{\bar{\sigma}_a} = \bar{\sigma}_a g_a(\bar{\sigma}_a)$ and therefore also disappears as $\bar{\sigma}_a$ gets large. Consequently, $\frac{\partial N_a}{\partial\bar{a}}$ tends to zero as $\bar{a}$ and $\bar{\sigma}_a$ — and budget size R — become large (for both L- and H-types).

By implication, as we keep increasing the budget size R, the left-hand side of (18) goes to zero, whereas the right-hand side increases without bound. As a result, for a large enough R the inequality has to be violated.

Part 3.

Suppose that the optimal policy satisfies $B = \bar{B}$. Because a full-separation policy delivers N_L^* individuals, an optimal policy must provide benefits to more than N_L^* individuals. By definition of $\bar{a}^*$, it must therefore satisfy $\bar{a} \geq \bar{a}^*$. Furthermore, it must satisfy $G_L(\bar{\sigma}_L) > N_L^*$ (the number of low-ability applicants which is greater than the number of recipients must be greater than N_L^*), such that $\bar{\sigma}_L > G_L^{-1}(N^*)$. Recall the identity $\bar{\sigma}_H = \frac{\bar{a}-a_H}{\bar{a}-a_L}\bar{\sigma}_L$. This formula is increasing in $\bar{a}$ and therefore $\bar{\sigma}_H \geq \frac{\bar{a}^*-a_H}{\bar{a}^*-a_L}\bar{\sigma}_L > \frac{\bar{a}^*-a_H}{\bar{a}^*-a_L}G_L^{-1}(N_L^*)$. Note that this lower limit is strictly positive because $\bar{a}^* > a_H$. Now, note that we also have an upper bound for $\bar{\sigma}_H$: We need to have at least N_L^* low-ability recipients and, with the budget R, we can then have no more than $\frac{R}{B} - N_L^*$ high-ability recipients. Consequently, $N_H < \frac{R}{B} - N_L^*$ while $N_H > P(0)\,G_H(\bar{\sigma}_H)$ (at least a share $P(0)$ of the high-ability applicants receive benefits, because $\tilde{P}(\alpha,B) > P(0)$). Consequently, $\bar{\sigma}_H < G_H^{-1}\left(\frac{1}{P(0)}\left(\frac{R}{B} - N_L^*\right)\right)$.

Putting it together we have $\frac{\bar{a}^*-a_H}{\bar{a}^*-a_L}G_L^{-1}(N_L^*) < \bar{\sigma}_H < G_H^{-1}\left(\frac{1}{P(0)}\left(\frac{R}{B}-N^*\right)\right)$. If the upper bound is lower than the lower bound, there is no $\bar{\sigma}_H$ that satisfies this condition and therefore our original assumption that $B=\bar{B}$ is optimal must be false.[24] ■

B A Welfarist Approach to Complexity

In this section, we outline a welfarist approach to complexity and discuss how it affects the main insights presented above. We consider a government maximizing the welfare of low-ability individuals, i.e. it puts a zero welfare weight on high-ability individuals. This objective retains the spirit of the poverty-alleviation problem considered in the rest of the paper, but shifts the focus from income to utility. Because low-ability individuals differ with respect to the σ-parameter, we consider the ex-ante expected utility of a low-ability individual drawing σ from the distribution $G_L(\sigma)$. Hence, the social welfare function is given by

$$N_L^A \cdot u(a_L + B - f(\alpha)) + N_L^R \cdot u(a_L - f(\alpha)) + \left(\bar{N}_L - N_L^A - N_L^R\right) \cdot u(a_L), \tag{B.1}$$

where $N_L^A = N_L^A(\alpha, \bar{a}, B)$ is the number of accepted low-ability applicants, $N_L^R = N_L^R(\alpha, \bar{a}, B)$ is the number of rejected low-ability applicants, and $\bar{N}_L - N_L^A - N_L^R$ is the number of low-ability individuals who decide not to take up the benefit. We can alternatively write this objective as

$$N_L^A \cdot \left[u_L^A - \bar{u}_L\right] + N_L^R \cdot \left[u_L^R - \bar{u}_L\right], \tag{B.2}$$

where $u_L^A \equiv u(a_L + B - f(\alpha))$, $u_L^R \equiv u(a_L - f(\alpha))$, and $\bar{u}_L \equiv u(a_L)$. We have dropped the term $\bar{N}_L \cdot \bar{u}_L$, which is a constant that does not depend on policy parameters. In the above expression, we have that $u_L^A - \bar{u}_L$ is positive, while $u_L^R - \bar{u}_L$ is negative. This implies that the government objective is increasing in the number of low-ability recipients N_L^A as in the earlier formulation, and decreasing in the number of rejected low-ability applicants N_L^R as the worst possible outcome for a low-ability individual is to incur the complexity cost without getting the benefit. In other words, this objective favors policies that substitute take-up errors (type Ia) for rejection errors (type Ib). This is in contrast to the poverty-alleviation objective, where the two forms of type I error are perfect substitutes. Finally, the utility terms $u_L^A - \bar{u}_L$ and $u_L^R - \bar{u}_L$ are obviously decreasing in the complexity costs, and the policy makers now account for this effect.

Will there be more or less complexity under the welfarist approach than under the income maintenance approach? There are two new effects. First and obvious, because

[24] Remarks:

1. There is no inconsistency with $B = \bar{B}$ for small R — as we reduce R, $a^* \to a^H$ and therefore the lower bound goes to zero (I think the upper bound also goes to zero, but apparently our assumption of a finite slope of density guarantees that it does not go to zero that fast).
2. If we can choose $B > \bar{B}$, a^* would fall — this is the same argument as the one we made to show that we can always spend all of the money on a full separation policy, higher B allows for setting higher α while holding $\tilde{P}$ constant. With the same $\tilde{P}$ but higher α, probability of receiving benefits for everyone goes up because $\bar{a} > a^L$ and screening is better — there is therefore more applicants and they are more successful. Consequently, for high enough B we can guarantee the existence of σ^H that would satisfy the inequality.

the utility cost of complexity enters the government objective directly, complexity becomes less attractive as an instrument. Second, the fact that rejection errors are now worse than take-up errors favors policies that are capable of identifying applicants more precisely even when this reduces the number of applicants. All else equal, this creates an argument for a more complex screening process.

To understand the trade-offs in program design, consider the maximization of social welfare (B.2) with respect to α, $\bar{a}$, and B subject to a budget constraint given by $\left[N_L^A + N_H^A\right] B \leq R$. An interior solution with $N_H^A > 0$ (non-separation equilibrium) has to satisfy the following first-order conditions for α and $\bar{a}$:

$$\frac{\partial N_L^A}{\partial \alpha}\left[u_L^A - \bar{u}_L\right] + \frac{\partial N_L^R}{\partial \alpha}\left[u_L^R - \bar{u}_L\right] + N_L^A \frac{\partial u_L^A}{\partial \alpha} + N_L^R \frac{\partial u_L^R}{\partial \alpha} = \lambda \left[\frac{\partial N_L^A}{\partial \alpha} + \frac{\partial N_H^A}{\partial \alpha}\right] B$$

$$\frac{\partial N_L^A}{\partial \bar{a}}\left[u_L^A - \bar{u}_L\right] + \frac{\partial N_L^R}{\partial \bar{a}}\left[u_L^R - \bar{u}_L\right] = \lambda \left[\frac{\partial N_L^A}{\partial \bar{a}} + \frac{\partial N_H^A}{\partial \bar{a}}\right] B$$

These two conditions can be combined to give

$$\frac{\partial N_L^A/\partial \alpha}{\partial N_H^A/\partial \alpha} = \frac{\partial N_L^A/\partial \bar{a}}{\partial N_H^A/\partial \bar{a}} - \left[\frac{\partial N_L^R/\partial \alpha}{\partial N_H^A/\partial \alpha} - \frac{\partial N_L^R/\partial \bar{a}}{\partial N_H^A/\partial \bar{a}}\right] \frac{u_L^R - \bar{u}_L}{u_L^A - \bar{u}_L - \lambda B} - \frac{N_L^A \frac{\partial u_L^A}{\partial \alpha} + N_L^R \frac{\partial u_L^R}{\partial \alpha}}{\frac{\partial N_H^A}{\partial \alpha}\left[u_L^A - \bar{u}_L - \lambda B\right]}$$

which is analogous to condition (14) in the poverty-alleviation model. The left-hand side term and the first right-hand side term correspond exactly to condition (14). These terms capture that the optimal setting of complexity and eligibility should reflect the capability of each instrument to screen low-ability recipients from high-ability recipients. The second and third terms on the right-hand side capture the two new effects mentioned above. The second term reflects that the instrument which is better at avoiding rejection errors should be pushed further. The third term reflects that, because complexity is associated with a direct utility cost, this instrument should be used less aggressively.

To conclude, intensity of testing and associated complexity are associated with the same screening benefits and costs as in the earlier formulation (operating through its effects on accepted applicants of the two types, N_L^A and N_H^A) along with an additional screening benefit from converting rejection errors into take-up errors (operating through N_L^R) and an additional negative effect from the utility cost of complexity (operating through u_L^A and u_L^R). Although equilibrium complexity may be lower in a welfarist model, the net effect on complexity is in general theoretically ambiguous. Moreover, because take-up errors are more acceptable than rejection errors under social welfare maximization, the model would still be able to sustain incomplete take up as an equilibrium policy outcome. In this sense, the model retains the central insights of the simpler poverty-alleviation model.

References

Aizer, Anna, "Public Health Insurance, Program Take-up and Child Health," *Review of Economics and Statistics*, August 2007, *89* (3), 400–415.

Akerlof, George A., "The Economics of "Tagging" as Applied to Optimal Income Tax, Welfare Programs, and Manpower Planning," *American Economic Review*, March 1978, *68* (1), 8–19.

Benítez-Silva, Hugo, Moshe Buchinsky, and John Rust, "How Large Are the Classification Errors in the Social Security Award Process?," Working Paper 10219, National Bureau of Economic Research January 2004.

Besley, Timothy and Stephen Coate, "Workfare versus Welfare Incentive Arguments for Work Requirements in Poverty-Alleviation Programs," *American Economic Review*, March 1992, *82* (1), 249–61.

_ **and** _ , "The Design of Income Maintenance Programmes," *Review of Economic Studies*, April 1995, *62* (2), 187–221.

Bitler, Marianne, Janet Currie, and John Karl Scholz, "WIC Participation and Eligibility," *Journal of Human Resources*, 2003, *38*, 1139–79.

Blackorby, Charles and David Donaldson, "Cash versus Kind, Self-Selection and Efficient Transfers," *American Economic Review*, September 1988, *78* (4), 691–700.

Boadway, Robin and Katherine Cuff, "Monitoring Job Search as an Instrument for Targeting Transfers.," *International Tax and Public Finance*, 1999, *6* (3), 317 – 337.

_ **, Nicolas Marceau, and Motohiro Sato**, "Agency and the Design of Welfare Systems.," *Journal of Public Economics*, 1999, *73* (1), 1 – 30.

Coady, David, Margaret Grosh, and John Hoddinott, *Targeting Transfers in Developing Countries: Review of Lessons and Experience*, Washington, DC: World Bank, 2004.

Cuff, Katherine, "Optimality of Workfare with Heterogeneous Preferences," *Canadian Journal of Economics*, February 2000, *33* (1), 149 – 174.

Currie, Janet, "The Take-Up of Social Benefits.," in Alan J. Auerbach, David Card, and John M. Quigley, eds., *Public Policy and the Income Distribution*, New York: Russell Sage Foundation, 2006, 80 – 148.

_ **and Jeffrey Grogger**, "Explaining Recent Declines in Food Stamp Program Participation," *Brookings-Wharton Papers on Urban Affairs*, 2001, 203–244.

Daly, Mary C. and Richard V. Burkhauser, "The Supplemental Security Income Program," in **Moffitt, Robert A., ed.**, *Means-Tested Transfer Programs in the United States*, University of Chicago Press.

Daponte, Beth Osborne, Seth Sanders, and Lowell Taylor, "Why Do Low-Income Households not Use Food Stamps? Evidence from an Experiment," *Journal of Human Resources*, Summer 1999, *34* (3), 612–628.

Diamond, Peter A. and Eytan Sheshinski, "Economic Aspects of Optimal Disability Benefits," *Journal of Public Economics*, May 1995, *57* (1), 1–23.

_ **and James A. Mirrlees**, "A Model of Social Insurance with Variable Retirement," *Journal of Public Economics*, 1978, *10* (3), 295–336.

Duclos, Jean-Yves, "Modelling the Take-Up of State Support.," *Journal of Public Economics*, 1995, *58* (3), 391–415.

Heckman, James J. and Jeffrey A. Smith, "The Determinants of Participation in a Social Program: Evidence from a Prototypical Job Training Program," *Journal of Labor Economics*, August 2004, *22* (2), 243–98.

Jacquet, Laurence and Bruno Van der Linden, "The Normative Analysis of Tagging Revisited: Dealing with Stigmatization.," *FinanzArchiv*, 2006, *62* (2), 168 – 198.

Kanbur, Ravi, "Transfers, Targeting and Poverty," *Economic Policy*, April 1987, *2* (4), 112–147.

_ **, Michael Keen, and Matti Tuomala**, "Labor Supply and Targeting in Poverty Alleviation Programs," *World Bank Economic Review*, May 1994, *8* (2), 191–211.

Keen, Michael, "Needs and Targeting," *Economic Journal*, 1992, *102* (410), 67 – 79.

Kleven, Henrik and Wojciech Kopczuk, "Transfer Program Complexity and the Take Up of Social Benefits," Working Paper 14301, National Bureau of Economic Research September 2008.

Kreiner, Claus Thustrup and Torben Tranæs, "Optimal Workfare with Voluntary and Involuntary Unemployment," *Scandinavian Journal of Economics*, 2005, *107* (3), 459 – 474.

Mill, John Stuart, *Principles of Political Economy*, Vol. Books IV and V, Harmondsworth, UK: Penguin, 1848.

Mirrlees, James A., "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, April 1971, *38* (114), 175–208.

Moffitt, Robert, "An Economic Model of Welfare Stigma," *American Economic Review*, December 1983, *73* (5), 1023–35.

Moffitt, Robert A., ed., *Means-Tested Transfer Programs in the United States*, University of Chicago Press, 2003.

Nagi, Saad Z., *Disability and Rehabilitation: Legal, Clinical, and Self-Concepts and Measurement*, Columbus, OH: Ohio State University Press, 1969.

Nichols, Albert L. and Richard Zeckhauser, "Targeting Transfers through Restrictions on Recipients," *American Economic Review*, May 1982, *72* (2), 372–77.

Parsons, Donald O., "Imperfect 'Tagging' in Social Insurance Programs," *Journal of Public Economics*, October 1996, *62* (1-2), 183–207.

Prendergast, Canice, "The Motivation and Bias of Bureaucrats.," *American Economic Review*, 2007, *97* (1), 180 – 196.

Ravallion, Martin, "How Relevant is Targeting to the Success of an Antipoverty Program," *World Bank Research Observer*, August 2009, *24* (2), 205–31.

Salanié, Bernard, "Optimal Demogrants with Imperfect Tagging.," *Economics Letters*, 2002, *75* (3), 319 – 324.

Slemrod, Joel, "Optimal Taxation and Optimal Tax Systems," *Journal of Economic Perspectives*, Winter 1990, *4* (1), 157–78.

_ **and Shlomo Yitzhaki**, "Tax Avoidance, Evasion and Administration," in Alan Auerbach and Martin S. Feldstein, eds., *Handbook of Public Economics*, Vol. 3, Amsterdam; New York: Elsevier/North Holland, 2002.

Stern, Nicholas, "Optimal Taxation with Errors in Administration," *Journal of Public Economics*, March 1982, *17* (2), 181–211.

FIGURE 1. DENSITIES OF P(.) AND CLASSIFICATION ERRORS

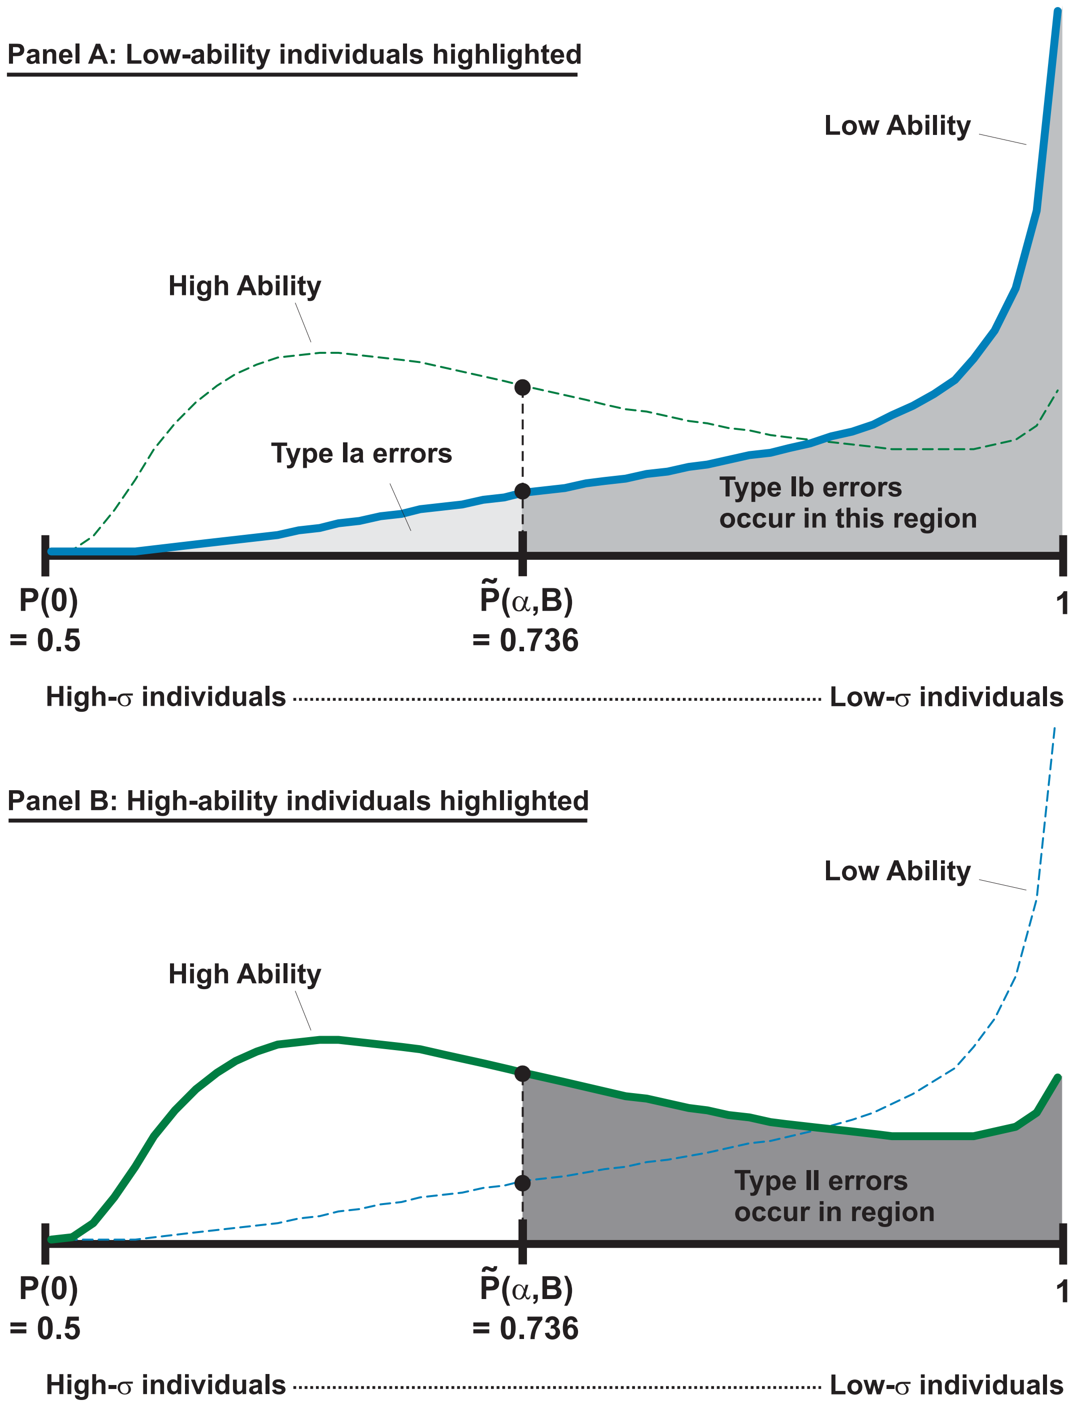

Parameters are as follows: $a_L = 1$, $a_H = 2$, $\bar{N}_L = 1000$, $\bar{N}_H = 1000$, $\bar{B} = 1$, $\beta = 2$, ε/σ is normal (so that $P(0) = \frac{1}{2}$), σ is assumed to be log-normal with a mean and variance equal to 1, $f(\alpha) = 0.5 \cdot \alpha$. The parameters of the program are selected optimally at the level of revenue equal to $R^* + 500 \approx 1310$, where R^* is the maximum budget allowing for the first-best allocation (see Section 3.4). The working paper version contains extensive discussion of simulation exercises that vary parameters of the model.