Games and Economic Behavior 104 (2017) 632–655

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Test design under voluntary participation

Frank Rosar *

Department of Economics, University of Bonn, Lennéstr. 37, 53113 Bonn, Germany

ARTICLE INFO

Article history:
Received 27 April 2016
Available online 17 June 2017

JEL classification:
D02
D82
D83

Keywords:
Test design
Bayesian learning
Concavification
False positive
Asymmetric information
Voluntary participation

ABSTRACT

An agent who is imperfectly informed about his binary quality can voluntarily participate in a test that generates a public signal. I study the design of the test that allows for optimal learning of the agent's quality when the agent strives for a high perception of his quality but is averse towards perception risk. For a large class of reduced-form utility functions that reflect these properties, the optimal test is binary and not subject to false positives. I uncover the forces that drive this result and show how the problem with endogenous participation can be transformed into a problem to that the concavification approach from the Bayesian persuasion literature applies. Furthermore, for a non-reduced version of my model where the designer estimates the agent's quality but suffers either more from false positives or from false negatives, I show that the same type of test is optimal.

1. Introduction

Technological advances have in many areas improved the capability to reveal information about an agent through testing procedures. In contrast to signaling and screening mechanisms that were subject of intense study in the past, such procedures are in many environments capable of revealing information that goes beyond what the agent knows himself. However, either because a regulating agency saw a need to protect the agent or for some technical reason, testing often requires the agreement of the agent. This article studies the design of the testing procedure in such environments.

To fix ideas, think of the following examples: A genetic test can determine whether a patient will develop a certain illness, but testing necessitates for legal reasons the formal agreement of the patient. Big data analysis can be used to assess whether a bank is viable in a scenario that is unknown to the bank, but the bank must grant access to the necessary data. Polygraphs and fMRI scanners can be used to judge about the eligibility of a potential employee, but using them requires the candidate's agreement. Similar problems arise also in more classical environments: The judge in an inquisitorial legal system controls the generation of information, but she requires an imperfectly informed prosecutor to bring cases to court.[1] Product testing can reveal information about the safety of a product, but the imperfectly informed producer has to apply for certification.[2]

In this kind of applications, the agent's quality is either good (i.e., he is healthy/viable/eligible/...) or bad (i.e., he is ill/nonviable/ineligible/...), but he is often only imperfectly informed about his true quality. A test is any device that, if

* Fax: +49 228 73 7940.
E-mail address: email@frankrosar.de.

[1] Kamenica and Gentzkow (2011) discuss in their leading example the complementary problem where information generation is controlled by an uninformed prosecutor.

[2] See Harbaugh and Rasmusen (forthcoming) for a deeper discussion of the certification application.

http://dx.doi.org/10.1016/j.geb.2017.06.002
0899-8256/

used, issues a public signal that depends directly on the agent's true quality.[3] An accurate test perfectly reveals the agent's quality. It generates a *positive* result if the agent is good and a *negative* result if he is bad.[4] Any binary test that is not accurate is subject to false positives, false negatives, or both. Similar interpretations can be employed for non-binary tests. After the test is designed by the principal, the agent decides upon participation. The principal uses the generated test result or the agent's decision not to participate to draw a Bayesian inference about the agent's quality. I consider a version of the model with stylized reduced-form utility specifications that depend directly on this inference. In addition, I discuss a non-reduced version where the principal has to take a decision and a class of utility functions with standard properties that is interesting from an applied perspective.

In general, the principal strives for uncovering information about the agent's quality. She dislikes false positives and false negatives. However, depending on the application, she may dislike one type of error more than the other. The agent cares about how his quality is perceived by the principal. He benefits when the perception is higher but he is in many problems averse towards perception risk. For instance, such an aversion may be due to the Hirshleifer effect or it may arise for psychological reasons.[5]

What makes the test design problem interesting is that there is no full unraveling under an accurate test. The agent faces a trade-off: non-participation signals unfavorable private information about his quality but participation comes along with a perception risk. The usual unraveling argument fails as the imperfection of the agent's private information limits how adverse the inference associated to non-participation can be. Because of his aversion towards perception risk, the agent with the worst possible private information strictly prefers to be perceived as having the worst possible private information to his quality being perfectly revealed by an accurate test. The relevance of this kind of effect was already observed by Stiglitz (1975):

> If individuals are very risk averse and not perfectly certain of their abilities, then they may prefer to be treated simply as average rather than to undertake the chance of being screened and labelled below average.

Thus, even if there are no technical constraints that limit the test accuracy and accuracy is costless, the principal may have an incentive to choose an inaccurate test to foster participation.

This gives rise to a number of questions: Is optimal learning achieved through an accurate test? If not, what is it optimal to test for? Is the optimal test subject to false positives, false negatives or both? How is this affected by whether the principal suffers more from false positives or from false negatives?

My article has two main contributions. The first contribution is theoretical. I use the version of my model with reduced-form utility specifications to uncover the forces that drive the optimal test design. For a large class of reduced-form utility functions that reflect the principal's interest in learning and perception risk-aversion on the agent's side, I find that a single test structure is optimal: the optimal test is binary and not subject to false positives.

In the literature on Bayesian persuasion (Kamenica and Gentzkow, 2011) the sender chooses a Bayes plausible distribution of beliefs to maximize her value function. Concavification methods can be used to solve this problem (see, e.g., Aumann and Maschler, 1995). Such an approach cannot be (directly) applied to my problem as two technical difficulties arise. First, the principal cannot completely control the distribution of beliefs about the agent's quality. This distribution is jointly determined by her test design and the agent's endogenous participation behavior. Second, as the agent is privately informed, two different distributions of posterior beliefs are relevant instead of one.[6] I show that it is possible to transform the problem in a way such that the concavification approach applies nevertheless.

The second contribution is more applied. I discuss the non-reduced version of my model where the principal has to estimate the agent's quality but suffers either more from a high estimate when it turns out that the agent is bad or from a low estimate when it turns out that he is good. That is, either false positives or false negative are more costly for the principal. Expectedly, tests that are not subject to false positives perform well when the principal suffers more from false positives; more surprisingly, such tests turn out to be also optimal when the principal suffers more from false negatives.

2. A model of test design with voluntary participation

2.1. The reduced model

A principal (she) wants to learn about the quality $\omega \in \Omega \equiv \{g, b\}$ of an agent (he). The agent's quality is either good ($\omega = g$) or bad ($\omega = b$), but the agent is only imperfectly informed about his quality. He knows that he is good with probability

[3] I am interested in the case where the generated information is always learned by the principal. Matthews and Postlewaite (1985) compare an agent's incentive to generate information through an accurate test for voluntary and mandatory disclosure of the generated information. See also Farhi et al. (2013).

[4] Note that I fix the wording such that "positive" refers to the result that is better for the agent. For some applications, the wording is interchanged. E.g., a positive HIV test is the result that is worse for the agent.

[5] According to Hirshleifer (1971), such an aversion can arise because the revelation of public information takes away hedging opportunities. Moreover, in medical contexts, learning aversion arises often for psychological reasons. An indicator for such an aversion is that participation rates in medical tests are often low even when testing is costless. See Lyter et al. (1987) and Hull et al. (1988).

[6] In this respect, my problem is related to the Bayesian persuasion problem with heterogeneous priors in Alonso and Câmara (2016a).