Towards an Economic Analysis of Trusted Systems
(Position Paper)

Dirk Bergemann[*]
Yale University
Economics Department
New Haven, CT 06520-8268
dirk.bergemann@yale.edu

Joan Feigenbaum[†]
Yale University
Computer Science Department
New Haven, CT 06520-8285
joan.feigenbaum@yale.edu

Scott Shenker[‡]
U. C. Berkeley and ICSI
EECS Department
Berkeley, CA 94720-1776
shenker@icsi.berkeley.edu

Jonathan M. Smith[§]
Univ. of Pennsylvania
CIS Department
Philadelphia, PA 19104
jms@cis.upenn.edu

February 27, 2004

1 Introduction

Trusted-platform initiatives such as Microsoft's Next-Generation Secure-Computing Base (NGSCB) and the industry-wide Trusted-Computing Group (TCG) effort are the subject of significant research and development now. The goal of these initiatives is to change a fundamental fact about networked, general-purpose computers that is often viewed as a barrier to security: Once data are sent from one machine to another, the sender loses control over them. Trusted-platform designs [AFS97] offer hardware-based, cryptographic support for proofs that a potential receiver's machine is running an approved software stack. By making such proofs prerequisites for the transfer of sensitive data, owners of these data can ensure that only authorized applications will be run and only authorized actions will be taken by users. The best publicized motivation for this type of "remote control" of networked

[*] Supported in part by NSF.

[†] Supported in part by ONR and NSF.

[‡] Supported in part by NSF.

[§] Currently on leave at DARPA. Smith's work on this paper was completed while he was at the Univ. of Pennsylvania.

computers is copyright enforcement for entertainment content, but in fact a wide variety of security-policy enforcement might be enabled if trusted platforms worked as advertised and were *widely adopted*.

In this position paper, we provide a simple model of a trusted-platform market, drawing upon recent work on the more general topic of two-sided markets. Using this model, we make some basic observations about adoptability of and alternatives to trusted platforms. Our main goal is to lay the groundwork for an economically informed research agenda on trusted platforms, and, toward this end, we suggest some specific open questions about how one could elaborate on our simple model.

2 A Basic Trusted-Platform Market Model

The introduction of a platform to a networked industry can significantly alter the market structure. Typically, the provider of the platform faces a *two-sided market*. For example, if a trusted platform for popular music is to succeed as a product, it must be adopted by the music-distribution industry, which is dominated by a few large firms, and also be accepted by popular-music consumers, a group that includes a significant fraction of the general public. The platform manufacturer thus has to court both sides of the market if widespread adoption is to happen. Two-sidedness of this nature is arguably a very common feature in markets with network externalities. Other prominent examples include markets for videogame consoles, browsers, streaming-media software, and operating systems.

Despite the pervasiveness of platforms and two-sidedness in markets with externalities, research about the economic functioning of these markets has flourished only recently (*e.g.*, [CJ01, CJ04, RT04]). We start with a static model of platform competition that shares many basic features of the models used in these works. For concreteness, we refer to "content providers" and "content consumers," but the model applies generally to two-sided markets with externalities.

There are a finite number of competing platforms, denoted by $i \in \{1, ..., I\}$. The platform market is assumed to have two sides: the content providers, denoted by $j \in \{1, .., J\}$, and the content consumers, denoted by $k \in \{1, ..., K\}$. The content provider j has to decide whether or not to adopt the platform offered by i. The decision by provider j regarding the platform i is denoted by $x_j^i \in \{0, 1\}$. Similarly, the content consumer k has to make an adoption decision with respect to platform i, and his decision is represented by $x_k^i \in \{0, 1\}$. We allow for the possibility of *multihoming*, *i.e.*, for a particular provider or consumer to have access to several platforms. Finally, the consumer has to make a decision about whether to purchase content from provider j on platform i, and his decision is denoted $x_k^{i,j} \in \{0, 1\}$.

Consumer k attaches a utility u_k^j to the content of provider j, and each seller j is assumed to offer only one product. The values of u_k^j are assumed to be random draws from a fixed and known distribution $u_k^j \sim F^j(\cdot)$, to account for consumer as well as provider heterogeneity. The value of a secure transaction for provider j is given by w_j, and the probability of the transaction's being secure on platform i is given by q_i.

The platform provider i can charge either side of the market a (one-time) registration fee (p_i^j and p_i^k, respectively) and/or a transaction (rental) fee, $r_i^j$and r_i^k, for provider j and consumer k, respectively. Separately, the content provider j charges the consumer k a transaction fee r_j^k for its content. For simplicity, the cost of providing the platform and the content is normalized to zero. The net values of the three groups of agents can then be expressed in terms of the registration fees $\mathbf{p} = \left\{p_i^j, p_i^k\right\}_{i,j,k}$, transaction fees $\mathbf{r} = \left\{r_i^j, r_i^k, r_j^k\right\}_{i,j,k}$, and finally platform-adoption and content-purchase decisions: $\mathbf{x} = \left\{x_j^i, x_k^i, x_k^{i,j}\right\}_{i,j,k}$. The revenue of platform provider i is given by:

$$V_i\left(\mathbf{p}, \mathbf{r}, \mathbf{x}\right) = \sum_{j=1}^{J} \left(\sum_{k=1}^{K} \left(x_j^i x_k^i x_k^{i,j} \left(r_i^j + r_i^k \right) + x_k^i p_i^k \right) + x_j^i p_i^j \right),$$

whereas the revenue of the content provider j is given by:

$$V_j\left(\mathbf{p}, \mathbf{r}, \mathbf{x}\right) = \sum_{i=1}^{I} \left(\sum_{k=1}^{K} x_j^i x_k^i x_k^{i,j} \left(q^i w_j + r_j^k - r_i^j \right) - x_j^i p_i^j \right),$$

and the net utility of content consumer k is given by:

$$V_k\left(\mathbf{p}, \mathbf{r}, \mathbf{x}\right) = \sum_{i=1}^{I} \left(\sum_{j=1}^{J} x_j^i x_k^i x_k^{i,j} \left(u_k^j - r_j^k - r_i^k \right) - x_j^i p_i^k \right).$$

The term $x_j^i x_k^i x_k^{i,j}$ indicates that, for a transaction between j and k to occur on platform i, it has to be the case that both provider and consumer have access to the platform i, or $x_j^i = x_k^i = 1$, and that consumer k decides to make a purchase on platform i from provider j, or $x_k^{i,j} = 1$. The current notation allows for perfect price discrimination, because the price of the platform can depend on the identity of the content provider j or the consumer k; however, as a baseline result, we are interested in a uniform price equilibrium, where, for all j, j', $p_i^j = p_i^{j'}$ and, more generally, where the registration or transaction fee does not carry a superscript and hence does not discriminate across providers or consumers. The value functions immediately display the presence of network externalities in this model, because the value of adoption (for either provider or consumer) depends on how many consumers or providers have adopted a particular platform. The model thus captures the essential aspects of network externalities and the presence of multi-product, multi-market contact.

3 Open Questions about Trusted-System Adoption

3.1 Market Structure and Governance

Suppose initially that there is only one privately owned trusted platform available. In this scenario, the only adoption decision for both providers and consumers is the choice between using the platform and not doing so. It is then immediate that the platform provider has substantial monopoly power that allows him to extract much of the surplus in the relationship between content providers and consumers. Because this tends to distort the efficiency of the equilibrium allocation, one is led to examine the nature of the inefficiency and potential remedies.

The platform acts as an intermediary between provider and consumer and may thus shift surplus from one group to the other. It is natural to ask whether the adoption of a trusted platform improves the welfare of one side of the market *vis-a-vis* the other side. This question is particularly relevant for entertainment-content providers, which are arguably few in number but large, and entertainment-content consumers, of whom there are many, each with (almost) zero market power.

The allocation resulting from a for-profit monopolist should be compared with the allocation resulting from a non-for-profit platform provider (such as Linux) or industry-sponsored platform (such as the current DVD standards).

Question 1 *What is the effect on the basic model of the number I of competing platform providers and of their governance structures (e.g., for-profit individual firm, not-for-profit individual firm, or industry consortium)?*

3.2 Time and Uncertainty

The basic static model of Section 2 conceives the adoption and purchase decisions as simultaneous and occurring under conditions in which all agents have complete information. Clearly, a dynamic and evolving description of the world would be more realistic. A dynamic version of the model would incorporate sequentiality of the adoption decision and determine the conditions under which a content provider or consumer would be the first one to adopt a new platform.

Question 2 *How do adoption decisions by content providers and consumers depend on the distribution of consumers' valuation of the different products, e.g., are high-value content providers more likely to adopt trusted platforms earlier?*

As with other software and hardware innovations, the reliability and safety of a "trusted platform" is initially uncertain, and thus the evolution of the perceived quality of a platform will influence the adoption and pricing decisions.

3.3 Multihoming and Interconnectivity

The existence of competing platforms and multihoming would seem to strengthen the competition among the platform providers and thus increase the efficiency in the market. On the other hand, there is often a fixed roll-out cost for a new platforms, and thus multihoming may lead to inefficient duplications.

Question 3 *What are the trade-offs presented by the possibility of multihoming? What incentives do competing platform providers have to allow interconnectivity?*

3.4 Information Asymmetry

Finally, platform providers and content consumers typically have very different information about the security vulnerabilities of a trusted platform. From basic microeconomic theory, we know that asymmetry in information can lead to inefficiencies in the adoption of and trading on a platform.

Question 4 *Under which conditions does a platform provider have incentives to disclose security vulnerabilities and other initially proprietary information about the trusted platform?*

Related work has recently been done by Schechter and Smith [SS03].

4 Alternatives to Trusted Platforms

We now discuss alternatives to trusted-platform computing, continuing to use entertainment content and DRM as an example. For simplicity, assume that the consumer has three basic alternatives for each product. He can use it only as authorized by the provider (as trusted platforms are designed to ensure), he can use it privately for a variety of different purposes (some of which may not be authorized), or he can make the product publicly accessible on the network. We attach to each of these choices (authorized, private, or network use) by the agent a utility level, denoted by u_a, u_p, and u_n respectively, with the following ranking:

$$0 < u_a < u_p < u_n < \infty.$$

Next, consider the possibility of a network-wide monitoring system that can stochastically observe communication among agents. To be more specific, suppose that unauthorized action is observed with probability q. Following an observation, the content provider can ask the user to adopt a trusted platform to prevent any further unauthorized use. For simplicity, assume that the consumer uses one instance of the product in every period, that future utility is discounted with discount factor $\delta \in (0, 1)$, and that every consumer has an infinite lifespan. In this context, the factor δ can also be interpreted as the probability of a future interaction between content provider and consumer. We can interpret δ as the frequency of interaction between consumer and content provider.

Because flexible private use of the content generates a higher utility for the consumer but could lead to network-wide use, there is a trade-off for the content provider between imposing a trusted platform on consumers and relying on a network-monitoring system. The trade-off can be resolved if his pricing policy can simulateously maintain the participation constraint:

$$u_p - r_p \geq u_a - r_a,$$

which induces the consumer to select the private use option at the respective price r_p and the incentive constraint:

$$(u_p - r_p) + \frac{\delta}{1-\delta}(u_p - r_p) \geq (u_n - r_p) + \delta\left(qV_a + (1-q)V_p\right),$$

where

$$V_a = \frac{1}{1-\delta}\left(u_a - r_a\right),$$

and

$$V_p = \frac{1}{1-\delta}\left(u_p - r_p\right),$$

are the sums of discounted future utility from authorized and private use of entertainment content, respectively. The incentive constraint states that the consumer prefers to maintain the private-use option rather than using the content networkwide with a risk of discovery (of probability q) and a resulting restriction to authorized use. The price difference between the "insecure version" of the product and the trusted-platform version, given by:

$$r_p - r_a = \frac{\left(u_p - u_n\right)\left(1-\delta\right) + \delta q\left(u_p - u_a\right)}{\delta q},$$

will then be positive if

$$\left(u_p - u_n\right)\left(1-\delta\right) + \delta q\left(u_p - u_a\right) > 0.$$

Because the first term $(u_p - u_n)$ is negative, and the second term $(u_p - u_a)$ is positive, the difference will be positive if either the frequency δ of interaction is sufficiently high or the probability q of detection is sufficiently high.

This outline of an argument shows that the content provider might be better off relying on an imperfect and stochastic network-monitoring technology than imposing a trusted platform. This will be the case whenever the value from flexible private use (and possibly minor copyright infringement) can be granted to the consumer, but network-wide violation can be detected frequently enough to discourage larger-scale and more serious copyright infringement.

Several research directions immediately suggest themselves, including:

Question 5 *Apply this model to the analysis of existing content-distribution businesses that allow for flexible private use, such as iTunes.*

References

[AFS97] W. A. Arbaugh, D. J. Farber, and J. M. Smith, "A Secure and Reliable Bootstrap Architecture," in *Proceedings of the 1997 IEEE Symposium on Security and Privacy*, pp. 65–71.

[CJ01] B. Caillaud and B. Jullien, "Competing Cybermediaries," *European Economic Review Papers and Proceedings* **45** (2001), pp. 797–808.

[CJ04] B. Caillaud and B. Jullien, "Chicken and Egg: Competition among Intermediation Service Prociders," to appear in *Rand Journal of Economics*, 2004.

[RT04] J. Rochet and J. Tirole, "Platform Competition in Two-Sided Markets," to appear in *Journal of the European Economic Association*, 2004.

[SS03] S. Schechter and M. Smith, "How Much Security is Enough to Stop a Thief?: The Economics of Outsider Theft via Computer Systems and Networks," in *Proceedings of the 2003 International Financial Cryptography Conference*, pp. 122–137.