

Trusting Privacy in the Cloud

Prüfer, J.O.

Document version:

Early version, also known as pre-print

Publication date:

2014

[Link to publication](#)

Citation for published version (APA):

Prüfer, J. O. (2014). *Trusting Privacy in the Cloud*. (TILEC Discussion Paper; Vol. 2014-047). Tilburg: Tilburg.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright, please contact us providing details, and we will remove access to the work immediately and investigate your claim.



Discussion paper

TRUSTING PRIVACY IN THE CLOUD

by
Jens Prüfer

December 1, 2014

TILEC Discussion Paper No. 2014-047
CentER Discussion Paper No. 2014-073

ISSN 2213-9532
ISSN 2213-9419
<http://ssrn.com/abstract=2533819>

Trusting Privacy in the Cloud *

Jens Prüfer

Tilburg University[†]

December 1, 2014

Abstract

Cloud computing technologies have the potential to increase innovation and economic growth considerably. But many users worry that data in the cloud can be accessed by others, thereby damaging the data owner. Consequently, they do not use cloud technologies up to the efficient level. I design an institution that attenuates this problem. The scheme is built around a private, nonprofit organization called *cloud association*, which is governed by representatives of both cloud service providers and users, and which sources the actual auditing and certification tasks out to independent certifiers. I show how this institution incentivizes providers to produce high data security, and users to trust them and pay a premium for their services. The cloud association simultaneously solves providers' adverse selection problem and certifiers' moral hazard problem. By credibly implementing certified/not certified decisions, it drastically reduces the technological complexity faced by users, which boosts trust in cloud services.

JEL classification: D02, D71, L24, L31, L86

Keywords: cloud computing; privacy; data security; trust; associations; certification; economic governance; private ordering; market design

*I am grateful to several seminar audiences at Tilburg University, at the A4Cloud General Meeting in Tilburg, the CPB in Den Haag, the 2014 meeting of the International Society for New Institutional Economics at Duke University, and especially to Daniele Catteddu, Sebastian Dengler, Shivaram Devarakonda, Joeri van Hugten, Eleni Kosta, Ronald Leenes, Maartje Niezen, Erin O'Hara O'Connor, Bastiaan Overvest, and Patricia Prüfer who provided valuable feedback on an earlier draft of this paper. The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no: 317550 (A4Cloud). All errors are my own.

[†]Tilburg School of Economics and Management, CentER, TILEC, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands; j.prufer@uvt.nl.

1 Introduction

Cloud computing is a model for enabling convenient, on demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction (US National Institute of Standards and Technologies). Typical cloud services are delivery of software, infrastructure, and storage over the Internet, based on user demand. The use of cloud computing technologies decreases the fixed costs and transaction costs of using computer power. It makes the supply of information technology (IT) resources more flexible for users and consequently reduces the risks of fixed, long-term investments in IT infrastructure. In short, cloud computing technologies have a huge upside potential increasing the efficiency of many IT applications and, thereby, innovation and economic growth (Etro, 2009). In 2014, global business spending for infrastructure and services related to the cloud is expected to reach an estimated \$174.2 billion, a 20 percent increase on 2013.¹

However, there are important impediments to the wide adoption of cloud computing. Many users—both individual consumers managing their personal files and businesses using IT services as input into their own production—worry that data outsourced to the cloud can be accessed by others, notably public authorities with legitimate or illegitimate objectives as well as legal and illegal private actors.² Such concerns regarding privacy, security, and data protection are a key obstacle that hinder the cloud industry to realize its full economic and technological potential (Cattedu and Hogben, 2009).³

In computer science, privacy and data security concerns are related to the concept of *accountability*, which refers to a situation, where both a cloud service provider and a user “should be able to check whether the cloud is running the service as agreed. If a problem appears, they should be able to determine which of them is responsible, and to prove the presence of the problem to a third party, such as an arbitrator or a judge” (Haeberlen, 2010); see also Pearson (2011). A high level of accountability refers to a high level of data security, which is not only implemented by a cloud service provider but also trusted by users. When a seller offers her services to users—at a certain announced data-security level—the seller has potentially already implemented a diverging true data security level. Therefore, the underlying economic problem is *adverse selection*, which stems from asymmetric information between the seller and the buyer about the true quality of the offered cloud service and which puts the entire market at the risk

¹<http://news.investors.com/technology-click/021414-690137-amzn-goog-msft-battling-for-growing-cloud-market.htm>.

²The recent hacking of iCloud pictures of celebrities has sparked such fears (<http://www.computerworld.com/article/2601905/apple-icloud-take-reputation-hits-after-photo-scandal.html>).

³In Europe, sales of cloud computing services trail those in the U.S. by at least two years, owing to concerns about privacy and fears that business secrets could be stolen in U.S. based cloud centers (New York Times, 2012).

of breakdown (Akerlof, 1970).⁴

One standard way to solve adverse selection problems is *signaling* (Spence, 1973): if a cloud service provider with high implemented accountability could credibly signal its quality to users, and if it was prohibitively costly for low accountability providers to mimic this signal, users could separate the wheat from the chaff and avoid the breakdown of the market for expensive high-accountability cloud services. Such a signal could be offered by handing out warranties, for instance, such that a user suffering from a data breach would have a legal right to ask the cloud service provider to reimburse her for damages originating from the data breach. If this strategy was implementable, only high-accountability providers would offer a warranty, a strategy that would be too costly for low-accountability providers, who would expect facing many warranty claims. Then, users could rationally trust the announcements of providers to offer high-accountability services, simply because cheating would not give the provider any advantage. The warranty would serve as an insurance policy.

Unfortunately, in cloud computing, it would be prohibitively costly for a service provider to offer a warranty that covers 100 percent of all data security breaches, as, according to industry representatives, complete data security does not exist in the cloud.⁵ Hence, handing out warranties would involve high expected payments for providers, which would drive up prices or even render all business models, including those involving high accountability, infeasible. Instead, as I will show below, the optimal accountability level—which is actually paid by users in equilibrium and leaves providers nonnegative profit—will typically cover less than 100 percent of all risk and leave some residual risk to data security with the user. If a cloud provider’s warranty would only cover certain risks, however, in case of a privacy infringement it would be impossible for most users to identify whether the infringement would be covered by the warranty, or not, due to lack of technical knowledge. Moreover, even a technology-savvy user would suffer from the difficulty that it would be very expensive to prove in front of a public court that a certain data security breach or privacy infringement was due to too low accountability of the provider.

One alternative mechanism that may solve the information asymmetry problem in the cloud is private *certification* of accountability. Certification agencies have a long history of supporting markets suffering from adverse selection, both in IT-related services and beyond (Rao, 1994, Corbett, Montes-Sancho, and Kirsch, 2005). If successful, a certification agency would only award a certificate as “accountable cloud service provider” to providers who actually implemented high accountability levels, such that the certification status of a seller could be used as a proxy by users to make cloud service purchasing decisions. However, it has been documented

⁴Even in bilateral relationships between firms outsourcing their information security management to managed security service providers, detecting security incidents and the efforts of providers to prevent them is difficult (Cezar, Cavusoglu, and Raghunathan, 2014). In the cloud, such issues are multiplied.

⁵Both the industry partners and other computer scientists involved in the EU A4Cloud project confirmed this view (<http://www.a4cloud.eu/consortium>).

by empirical research that certified sellers are sometimes even less trustworthy than uncertified sellers (Edelman, 2011). Theoretical research has shown that certification agencies can be motivated by profit-maximization motives to only award certificates truthfully (Strausz, 2005). Such truthful behavior depends on the quick and reliable distribution of information about a provider’s true accountability level—and thereby about a certifier’s true efforts in identifying that accountability level and its truthful certification decisions—to all other users. Only then could users identify “captured” certification agencies, not trust their certificates anymore, and “punish” transgressors by not buying their services in the future. Given that the cloud computing industry is characterized by users with heterogenous demands for cloud services and heterogenous technological competence, and that the sheer amount of sellers and buyers on this global market makes it unlikely that the relevant information spreads to a sufficient share of users quickly, the *moral hazard* problem of a certifier, to award certificates truthfully, is overwhelming.

These considerations lead to the following research questions. How can we incentivize cloud service providers, who claim to have produced high accountability (or data security) levels, to keep their contractual obligations and to install procedures that protect the accountability needs of users? How can we make sure that users trust the promise of providers to implement certain levels of accountability? How can we reduce the information costs for users, many of whom lack the relevant knowledge or means to evaluate the accountability level of a given provider, such that they can make informed consumption decisions? If the *adverse selection* problem between cloud service providers and users could be solved by introducing some type of certification scheme, how could the resulting *moral hazard* problem of the certifier, to save on the effort cost to understand a provider’s true accountability level or to be bribed by the provider, be solved?

In an attempt to answer these questions, I design an institution that attenuates the specific problems of the cloud computing industry stated above: a coherent set of behavioral and enforcement rules for the transactors involved such that nobody has an incentive to change her behavior and such that this behavior involves using highly accountable cloud computing services.⁶

In order to get there, in Section 2, I first discuss the relevant literature and identify the type and key characteristics of the optimal institution. In Section 3, I construct and analyze a game-theoretic model of the market interaction between cloud service providers and users (and third parties supporting their transaction). That model allows to study the incentives of each party involved, taking the incentives of the other parties into account, and to characterize

⁶Although this mechanism has some general properties and can inspire other applications (see the conclusion), it is important that it is tailored towards the specific circumstances of the cloud computing industry. Otherwise, trust can be lost too easily. See Montiel, Husted, and Christmann (2012) for a good example how sensitive to the environment trust in private certification standards is.

the optimal and the equilibrium levels of accountability. The theoretical results are instructive both for managers of firms selling or buying cloud services and for policy makers concerned with Internet governance. Section 4 contains extensions and robustness checks, whereas section 5 discusses managerial and policy implications and concludes. The model analysis is relegated to the appendix.

2 Institutional choice: Reputation, litigation, or certification?

Under which conditions can users trust providers' announced accountability levels—and, thereby, privacy and data security promises? In search of the optimal institution that may solve this problem,⁷ a few key factors can be identified. First, due to the international character of the cloud computing industry, any single national or regional legislation or regulation is incapable to set and enforce behavioral rules because providers could just move their cloud resources to countries with less restrictions (possibly even without users noticing).⁸

Second, as exemplified by a 20 percent revenue growth rate in 2014 (cf. footnote 1), the cloud computing industry is highly innovative and dynamic. Consequently, technological possibilities constantly change, which requires continuous adaption of products, services, and business models. The administrative processes in bureaucratic, public administrations, and legislative and regulatory agencies, which often require many decision-makers to get informed and vote about topics at stake or to seek approval of new initiatives by elected government representatives, are not well suited to create and constantly adjust appropriate rules that reflect the technical state of the cloud computing industry.

Moreover, in many countries, data and privacy protection is a task of the same governmental departments that are in charge of public safety, namely ministries of the interior. This creates a moral hazard problem for these governmental departments because public safety is more effective with access to more (personal) data of citizens, whereas the goal of data protection is precisely to limit the flow of such data. Bundling both tasks within the same governmental department is therefore likely to produce “compromises” between these goals, which diminishes the to-be-expected protection of data that both individual end users and business users expose to the cloud.⁹

As a consequence of these concerns, I conclude that the optimal institution to govern the cloud is not characterized by public ordering but by private ordering, also referred to as self-

⁷Masten and Prüfer (2011) offer a classification of available governance institutions.

⁸Current struggles about the future governance of web-based industries among the multitude of stakeholders surrounding ICANN and the Internet Governance Forum provide evidence for the difficulty to find one single institution governing global industries (EU Commission, 2014). By contrast, the solution suggested here can start at small scale and flexibly grow over time.

⁹Prüfer (2013) presents these arguments in more detail.

governance, which is subject to more flexible rules and decision-making procedures.¹⁰

Amongst private ordering institutions, reputation mechanisms such as online feedback systems on exchange platforms, which offer transactors to leave feedback about their trade partner and hence create a reputation for honesty or high quality, have been successfully used in online environments (Dellarocas, 2003, Aperjis and Johari, 2010). But pure information feedback mechanisms may be unreliable and suffer from omitted or biased feedback (Dellarocas and Wood, 2008, Cabral and Hortacsu, 2010). Moreover, as long as the incentives and legal obligations of the party managing the exchange platform are unclear to users, users may not trust that they obtain complete and accurate information about the trading history of a given seller.¹¹

Theoretical research has shown that both decentralized institutions, such as social networks, and centralized institutions without an effective enforcement mechanism, such as trading platforms, only work well as long as the importance of traders is rather symmetric, and if the total number of traders is limited (Greif, Milgrom, and Weingast, 1994, Dixit, 2003). Empirically, Ostrom (1990) confirms these results in several case studies, and Christmann and Taylor (2006) find that ISO-certified firms in China set their level of compliance with the certification standard strategically, depending on expected sanctions and customer monitoring.

If the alternative to a reputation mechanism is litigation, reputation performs badly if the cost of litigation is low or the ability of courts to identify actual behavior of traders is high (Bakos and Dellarocas, 2011, Masten and Prüfer, 2014). Both requirements are not met in cloud computing, however, thereby diminishing the role of litigation in supporting contract enforcement in the cloud, on top of the other problems of public ordering discussed above. Bakos and Dellarocas (2011) show that litigation is more effective in preventing moral hazard (that is, seller shirking) but that reputation is better in avoiding adverse selection (if there are high quality and low quality sellers). Cai et al. (2013) obtain a related result by studying trading data from a Chinese online platform. They confirm that more buyer protection can worsen the adverse selection problem because it reduces the number of honest sellers, which, in turn, decreases buyers' trust that sellers do not act opportunistically. As an alternative, law-based approach, O'Hara (2005) explores how contract law could be used to increase trust of consumers in unknown online vendors. She concludes (p.1885): "Even if harmonization of the online consumer protections could be achieved, however, it is not clear that consumer trust in online vendors would be significantly enhanced."

In this paper I show that, even in an industry where litigation is too costly and too inflexible, and where meeting the courts' standard of proof is very hard, pure reputation-based information

¹⁰See Williamson (2002) for a definition and discussion of the concepts of public and private ordering.

¹¹On top, the policies of some governments to reserve the right to access data stored by firms registered in their countries under certain circumstances, shatters the trust of users that access to data stored in the cloud is limited to those parties defined by the data owner (Soghoian, 2010).

distribution mechanisms, such as those used on many trading platforms on the Internet, can still be less effective and less efficient than mechanisms that comprise active investigation of the facts and punishment by a central (private) party. Therefore, the solution offered to the question which type of institution optimally supports contract enforcement and prevents opportunistic behavior by cloud service providers is given by *certification agencies*: private organizations with an own legal entity that are staffed by industry experts, audit and monitor cloud service providers, and award trust certificates only to those providers whose accountability standards meet certain criteria.¹²

A first advantage of certification schemes is given by their reduction of information complexity. In a market with relatively few technology-savvy users and with the possibility of providers' accountability measures to differ in many dimensions, a certification scheme can pre-define minimum thresholds for each dimension—for instance, by specifying security standards for software, hardware, and physical access to data centers. Consequently, a multidimensional performance vector of cloud service providers is mapped onto a binomial, easily observable outcome variable: is a provider certified, or not? Buehler and Schuett (2014) show by means of a game-theoretic model that certification mechanisms outcompete minimum quality standards when the share of uninformed consumers is large. This condition is given in cloud computing, where few users are IT experts (at the level required to judge the security of cloud services) and most users are unconnected to each other.

Apart from facilitating information flows about a provider's accountability level, the critical open issue is the certification agency's own *credibility problem* (Gibbons and Henderson, 2012). How to ensure that the certifier is not captured herself and that she does not have any incentive to misreport the findings of her audit, for instance, because of bribes received from the certified provider?¹³

This question—if not in the context of cloud computing—has already been studied in several theoretical contributions (Strausz, 2005, Mathis et al., 2009, Peyrache and Quesada, 2011). These authors show that reputation concerns of profit-maximizing certification agencies can mitigate their incentives to shirk and avoid that they award certificates to providers with insufficient quality (or data security) standards. The main line of argumentation is that, if a certifier would award a false certificate, consumers buying the falsely certified product would recognize

¹²Rao (1994) offers early empirical evidence that winning a certification contest enables firms to acquire a reputation for competence. Bae et al. (2013) show that firms wishing to secure an S&P rating react strategically by reporting more conservative financial statements. Corbett, Montes-Sancho, and Kirsch (2005) find in the context of ISO 9000 certification that certified firms make significant abnormal profits after certification. Certification even has measurable positive effects on the macro level: ISO 9000 certification levels are positively correlated with a country's export growth (Potoski and Prakash, 2009).

¹³Edelmann (2011) provides striking evidence that, without proper scrutiny in the auditing process, web sites with "trust" labels are *more* likely to behave untrustworthy than uncertified sites.

its low quality and, hence, conclude that the certifier was “captured” (Strausz, 2005). By playing a trigger strategy, that is, ceasing to buy from providers with certificates from captured certifiers, consumers could destroy future profits of those certifiers, which mitigates certifiers’ incentives to get captured in the first place. Studying the business of credit rating agencies, Mathis et al. (2009) conclude that reputation concerns can mitigate shirking if a sufficiently large fraction of a rating agency’s income is generated by other sources than rating financial products, for instance, by consulting. Peyrache and Quesada (2011) show that the probability that a certifier colludes with her client depends on the certifier’s pricing strategy—a problem that the institution described below takes care of by removing authority over pricing from the certifier.

However, a decentralized economic governance scheme that rests only on the reputation concerns of profit-maximizing certification agencies offering to audit providers based on idiosyncratic standards is unlikely to be effective in the cloud computing industry.¹⁴ The main reason is that such a scheme depends on the quick, cheap, and reliable transmission of information about a given certifier’s behavior (and hence her reputation) among all industry participants. This condition is hardly met in cloud computing, where heterogeneous users from all over the world, many of which do not have access to IT experts, cannot reliably judge whether a certificate is credible, or not. Moreover, if certifiers are directly paid by providers—just as rating agencies are paid by banks issuing financial products or certifiers of socially beneficial forest management are paid by timber producers,¹⁵ and as is argued for by Stahl and Strausz (2011)—they cannot be expected to value users’ interests more than necessary.

Strikingly, the theoretical literature modeling certification procedures has so far only considered one-layered certification schemes, where besides consumers and producers there is one or a few competing certification agencies. When reviewing certification schemes that are used in practice, however, it becomes prevalent that most schemes comprise two layers of certifiers. For instance, the ISO 27001 Framework, which specifies the requirements for managing a documented information security management system, assumes one layer of certification bodies and a second layer of national accreditation bodies, who certify the certifiers. The same holds for audit frameworks set up by Online Trust Services, the German Federal Office for Information Security (BSI), the UK Communications Electronics Security Group (CESG), the US Federal Information Security Management Act (FISMA), and nearly all other frameworks studied by ENISA (2013), which covers banking, payment cards, electricity generation, and health care providers, on top of several existing IT auditing schemes.

Therefore, in this paper, I build on the insight by practitioners that a two-layered certification scheme may be able to solve the problem of captured certifiers. This scheme is built

¹⁴Dixit (2009) offers a definition and overview of the concept of economic governance.

¹⁵See <https://ic.fsc.org/index.htm>.

around a private nonprofit organization that I call *cloud association*, which is governed by both representatives of providers and users and which sources the actual auditing and certification tasks out to a pool of independent for-profit certifiers.¹⁶

I show how and under which conditions such an institution can induce an equilibrium where cloud service providers produce high accountability levels and users trust them and buy their services, for a premium. In this equilibrium, a provider chooses to produce a high accountability level, which meets the certification requirements determined by the cloud association, because it increases her profits. It increases her profits because users are willing to pay a premium for a provider's services who is certified by the association. The certificate is reliable because certifiers have an incentive to invest effort in auditing and to honestly decide about the certification status of providers. This incentive exists because users who suffer from low accountability of a certified provider have an incentive to complain with the association. Following a complaint, the cloud association would investigate the case and, if it finds a breach of contract, revoke the provider's certificate and ban the captured certifier from all future business. Cheated users would actually complain because they are reimbursed for damage suffered from low accountability if the association finds for them. In turn, the association would keep its promises because of the checks and balances institutionalized in the governance structure of the association's board.

In the terminology of Bakos and Dellarocas (2011), the main problem between users and cloud service providers, once they implemented a certain accountability level, is adverse selection: providers with low accountability have an incentive to mimic those with high accountability, which can lead to market breakdown due to asymmetric information. This problem is solved by (private) litigation, in the sense that providers have to earn a certificate from a certification agency in order to get access to users with high willingness-to-pay. The higher-level problem, moral hazard of certification agencies who would like to save on the effort cost when auditing providers, is solved by a combination of private litigation (by the cloud association board's tribunal) and reputation (by withdrawing profitable future business from transgressing certifiers). Finally, the risk of capturing the cloud association's board is prevented because representatives of providers and users have to make decisions jointly. This governance structure implements cooperative behavior between the two opposing market sides, which makes sure that nobody is worse off than without it.

3 A Model of Cloud Governance

On the demand side of the market, consider a unit mass of risk-neutral users—who can be business or individual consumers—each demanding one unit of cloud computing services. Without

¹⁶Business associations can perform many other tasks that are valuable for their members. See Prüfer (2014) for details.

consumption, a user gets zero value. User i obtains indirect consumption utility of:¹⁷

$$v_i = u - (1 - \sigma)s_i - p, \quad (1)$$

where $u > 0$ denotes the gross utility from consuming a secure service, and $\sigma \in [0, 1]$ denotes the probability that no data security incident occurs. σ is a measure for the accountability level set by the provider selling the cloud service, which is unobservable to users.¹⁸ $s_i \sim U(0, \bar{s})$ denotes the disutility of user i in case of a security incident, which captures heterogeneous preferences for accountability across users.¹⁹ Service providers know the distribution of s_i but do not know the preferences of a specific user. Hence they cannot engage in perfect price discrimination. Higher \bar{s} refers to a more security-sensitive set of users on average. p stands for the price that a provider charges for one unit of her cloud services. We only consider cases where the following assumption holds.²⁰

Assumption 1 (User preferences) *Some users have strong accountability preferences: $\bar{s} > \frac{u}{2}$.*

On the supply side of the market, we focus on one monopolistic cloud service provider.²¹ The provider maximizes profits and solves:

$$\max_{\sigma, p} \pi = pq(p, \tilde{\sigma}) - c(\sigma), \quad (2)$$

where q denotes demand for the provider's services, which negatively depends on price p and positively on the accountability level expected by users, $\tilde{\sigma}$. To produce services of accountability σ , the provider incurs a cost $c(\sigma)$, which is increasing and convex in σ :

$$c(\sigma = 0) = c'(\sigma = 0) = 0, c'(\sigma > 0) > 0, c''(\sigma) > 0. \quad (3)$$

All other costs of the provider are irrelevant for this study and therefore normalized to zero.

Aiding the market transaction between users and providers, there are two types of intermediaries. First, there is a private, nonprofit organization called *cloud association*. The provider

¹⁷See section 4 for a discussion of the utility function.

¹⁸In practice, accountability is a multifaceted concept (Pearson et al., 2012). The service level agreement between a provider and a user typically describes how the provider promises to protect the user's data. I assume that users can distinguish more *promised* protection from less *promised* protection of their data. They do not observe the *realized* protection of their data, however.

¹⁹Both Chellappa and Sin (2005) and Solove (2007) provide evidence for the fact that users have very diverse privacy preferences and that virtually everyone is concerned about privacy infringements to some degree. See the literature cited therein for legal and philosophical discussions about the conception of privacy. Goldfarb and Tucker (2012) provide evidence for the fact that users' privacy concerns are increasing over time.

²⁰Tsai et al. (2011) show in experimental research that users are willing to pay a premium to online retailers who better protect their privacy. Pearson et al. (2012) discuss heterogeneous consumer preferences for accountability in the cloud.

²¹Section 4 discusses competition among several providers (and certifiers).

can become a member of this association for a fee F , which, in exchange, allows her to demand that her cloud services are audited and, if found to be of a certain accountability standard, that she is certified as “accountable cloud service provider” by the cloud association. Specifically, the certificate should be awarded if, and only if, the provider’s $\sigma \geq \hat{\sigma}$, where $\hat{\sigma}$ is the publicly known threshold accountability level determined by the cloud association.

The cloud association is governed by a board, which consists of representatives of both providers and users. For simplicity, I assume that the representatives of providers have formal authority over decision making but that users’ representatives can veto them.²² Hence, all decisions issued by the association’s board must be acceptable to both users and providers, by construction.

As the cloud association is a nonprofit organization, board members cannot legally take out profits (if existing) and can therefore be expected to pursue the objectives for which they were selected—advocating users’ (1) or providers’ (2) interests, respectively.²³ These control rights and internal checks and balances make the cloud association board the final, credible, and legitimate authority of the cloud industry, “which depends on acceptance by the cloud computing community” (Reed, 2013:5).²⁴

The association’s board, however, is limited in size and capacity and, therefore, cannot audit and certify providers itself. Therefore, the auditing procedure is outsourced to a profit-maximizing certification agency (or auditor).²⁵ An auditor can decide whether to thoroughly scrutinize the accountability level of a provider ($a = 1$) or only to pretend that she is truly auditing ($a = 0$). a is unobservable to everybody else, besides the auditor and the audited service provider, and comes at cost ac_A to the auditor, where $c_A > 0$. If the auditor chooses $a = 1$, she perfectly learns the provider’s accountability level, σ . If $a = 0$, the auditor does not learn anything.²⁶ In a second decision, which is costless, the auditor decides whether to award the provider with a certificate, or not.

To avoid that auditors are captured or bribed by the service provider they are to audit (Strausz, 2005), auditors have to seek a license from the cloud association. In practice, seeking

²²If both groups had active decision making power, I would have to make an assumption about the distribution of bargaining power. The quality of the results is not affected by the assumption used but the solution is more tractable.

²³Filistrucchi and Prüfer (2014) provide theory and evidence that nonprofit board members make decisions in line with their personal, observable values, thereby supporting upper echelons theory (Hambrick and Mason, 1984).

²⁴Users could be represented, for instance, by consumer protection agencies. Reed (2013) also points to the fact that industry self regulation works imperfectly when only one side of the market, such as producers, is in charge of governing the rules and their enforcement, as happened in the UK financial services sector.

²⁵The certification process has to comprise unexpected and ongoing checks of the provider’s implemented accountability level. See Probst et al. (2012) for more details regarding the technical requirements.

²⁶Reality is more nuanced but this assumption captures the qualitative costs and benefits of effort in auditing.

a license implies that the auditor is being educated by the cloud association how to conduct an audit using the association’s auditing protocol. This comes at a cost c_L , borne by the association.²⁷ If an auditor undergoes the licensing procedure, it is possible for the cloud association to verify ex post whether the auditor actually complied with the auditing protocol when auditing a provider.²⁸

The association board determines the certification protocol (captured by $\hat{\sigma}$). It also licenses certifiers and determines the remuneration sum R that it pays a certification agency in exchange for auditing a provider. To deter providers and certifiers from cheating about the true level of σ , the cloud association applies several policies: First, it offers users who bought the service of a certified cloud service provider the opportunity to complain about the provider’s accountability level.²⁹ If a user complains that $\sigma < \hat{\sigma}$, the association automatically starts an investigation. For the cost c_V it verifies whether the auditing procedure applied by the certifier to the cloud provider was in line with the protocol (that is, whether $\sigma \geq \hat{\sigma}$ and hence whether the certificate was awarded correctly, or not). By assumption, fraud committed by licensed auditors can be perfectly detected.³⁰ If the cloud association finds that a certificate was granted despite $\sigma < \hat{\sigma}$, it revokes the certificate from the provider and the license from the fraudulent certifier and bans the certifier from ever seeking a license again. In this case, the user who correctly complained is reimbursed by getting the service of an alternative certified provider, paid for by the cloud association, at accountability level $\hat{\sigma}$. If the association finds, following a complaint by the provider, that the certificate was not granted despite $\sigma \geq \hat{\sigma}$, it also revokes the certifier’s license but directly awards the provider with a certificate.³¹ For tractability reasons, I assume that complaining is costly but not very much so.³²

Assumption 2 (Costs of complaining) *Complaining about a provider’s accountability level to the cloud association costs a user $c_C = \varepsilon$, where $\varepsilon \rightarrow 0$.*

²⁷Whether the cloud association pays c_L (and increases the membership-fee F) or the certification agency pays c_L (and gets increased per-audit revenue R , which increases the membership-fee F) is equivalent within this model framework.

²⁸Using its auditing protocol also enables the association to verify ex post, that is, after an incident occurred, whether the incident occurred because of low security measures ($\sigma < \hat{\sigma}$) or because of the residual risk uncovered by the Service Level Agreement between the provider and a user. In this model, the residual risk is $1 - \sigma$.

²⁹Because complaining is endogenous to every user, this mechanism picks up the idea that voluntary feedback from traders in reputation systems is important in creating trust in markets (Li and Xiao, 2014).

³⁰In the model, making the certification decision verifiable is the reason for spending the licensing cost c_L . In practice, having a standardized auditing protocol (including the results and log files of the auditor’s tests) makes it much easier for an auditor of auditors to identify irregularities and hence fraud.

³¹It is possible that a user suffers from a security incident despite the fact that the certified provider set $\sigma \geq \hat{\sigma}$. In this case, the ex post verification of the cloud association would show that there was no fraud. Hence, neither the provider nor the certification agency would be punished. The user would suffer from the materialized residual risk inherent in imperfect security technologies. See also Footnote 28.

³²This assumption is discussed and relaxed in Section 4.

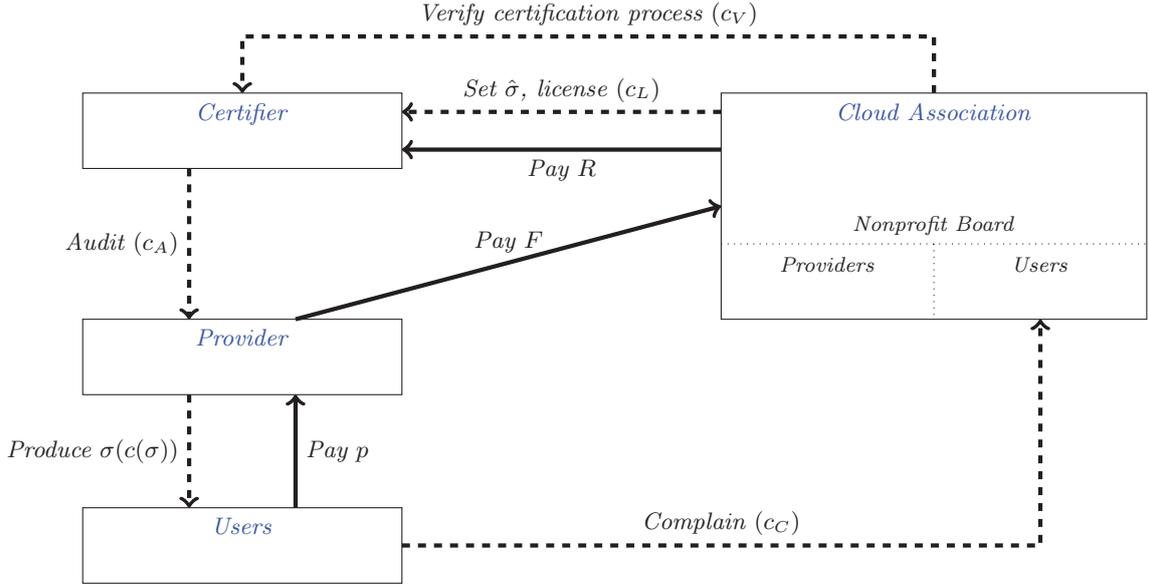


Figure 1: Governing the Cloud with the Cloud Association

An additional cheating prevention is the specification that there are no direct payment flows from the provider to her auditor (payments are channeled through the association) and that the cloud association sustains an entire pool of certification agencies, from which it randomly picks one certifier to audit the provider.³³

The proposed governance structure of the cloud computing industry is depicted in Figure 1. Payment flows are indicated by straight arrows, whereas all other services offered from one party to another one are indicated by dashed arrows (with their respective costs in parentheses).

The transactors play an infinitely repeated game with common discount factor $\delta \in (0, 1)$. The timing of the game in each period t is as follows:

1. The cloud association determines the membership-fee F and the accountability threshold level $\hat{\sigma}$, which it announces in public, and licenses a certifier (creating cost c_L).
2. The cloud service provider sets accountability level σ (for cost $c(\sigma)$) and price p , and decides whether to join the association and demand being audited (for membership-fee F).
3. If an audit is demanded, the association commissions a licensed certifier to audit the provider, offering revenue R .
4. The commissioned certifier chooses auditing effort a (for cost ac_A) and awards a certificate, or not. If the audited provider feels cheated by the certifier, she can complain with the cloud association.

³³This serves to minimize the opportunity for a specific certifier and the provider to engage in a repeated game, where payments in one period could be reciprocated by a positive certification decision in the future.

5. Users learn the provider's price and certification status and individually decide whether to buy the provider's service.
6. Every buyer can file a complaint with the association that a certified provider produced $\sigma < \hat{\sigma}$. All complaints are investigated by the cloud association (for cost c_V).

As every infinitely repeated game, this game has multiple equilibria. In Appendix A, I solve it for a subgame-perfect equilibrium in pure strategies that produces $\sigma > 0$ and prove the following Proposition.³⁴

Proposition 1 (Subgame-perfect equilibrium) *If all cost parameters ($c_C, c_A, c_L, c_V, c(\sigma)$) are not too high and utility u is not too low, then a subgame-perfect equilibrium exists and is characterized as follows:*

1. The cloud association sets a membership-fee $F^* = \frac{c_A}{\delta} + c_L + c_V$ and a threshold accountability level:

$$\hat{\sigma}^* = \begin{cases} \hat{\sigma}_1 & \text{if } \frac{u^2}{4(1-\hat{\sigma}_1)\bar{s}} - c(\hat{\sigma}_1) > u - (1 - \hat{\sigma}_2)\bar{s} - c(\hat{\sigma}_2), \\ \hat{\sigma}_2 & \text{otherwise} \end{cases} \quad (4)$$

where $\hat{\sigma}_1 \equiv \{1 - \frac{u}{2\sqrt{\bar{s}c'(\sigma)}} | \sigma < 1 - \frac{u}{2\bar{s}}\}$ and $\hat{\sigma}_2 \equiv \{\sigma | c'(\sigma) = \bar{s} \wedge \sigma \geq 1 - \frac{u}{2\bar{s}}\}$.

2. The provider produces $\sigma^* = \hat{\sigma}$, joins the association, demands certification, and asks a price:

$$p^* = \begin{cases} \frac{u}{2} & \text{if } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \\ u - (1 - \hat{\sigma})\bar{s} & \text{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}. \end{cases} \quad (5)$$

3. The association commissions a certifier and offers revenue $R^* = c_A/\delta$.
4. The certifier sets $a^* = 1$ and awards a certificate truthfully. The provider does not complain.
5. Users with accountability preferences $s_i \leq \bar{s}_1^*$ buy from the certified provider; users with preferences $s_i > \bar{s}_1^*$ do not buy. $\bar{s}_1^* = \frac{u}{2(1-\hat{\sigma}^*)}$ if $c'(\hat{\sigma}^*) < \bar{s}$, and $\bar{s}_1^* = \bar{s}$ if $c'(\hat{\sigma}^*) \geq \bar{s}$.
6. Because there is no fraud, nobody complains. But in case of fraud, a cheated user would complain about the provider to the cloud association.

Proposition 1 is a summary of Lemmas A.1 through A.5 (all in the appendix). It provides us with several results. First, the cloud association sets the membership fee for a provider, F^* , to the smallest possible level that can finance the system—simply because there is no use of

³⁴Equilibrium decisions are marked by asterisks.

piling up money within a nonprofit organization if the association members can just keep and spend the money individually as they like. Second, that “smallest level” just makes sure that the association has the means to pay for the licensing and, if necessary, also for the ex post verification of the certifier who is to audit the member. Additionally, the membership-fee has to cover the association’s expenses for paying the certifier an *efficiency wage*, which incentivizes her to invest effort in auditing and make a truthful certification decision.³⁵ This efficiency wage, R^* , grows in the cost the auditor incurs when exerting effort (c_A) and if the auditor is less patient (if δ decreases). Thereby, the auditor makes positive profits in equilibrium—that is, she earns an information rent—but behaves as aspired by the association because she is afraid of losing those profits in the future when her fraudulent behavior would be detected.

A second set of results concerns the market interaction following the association’s determination of the threshold accountability level, $\hat{\sigma}$, the corresponding incentives of the provider to seek certification, the provider’s pricing decision, and the users’ consumption decisions. The first insight here is that, although σ has support over the interval $[0, 1]$, as soon as $\hat{\sigma}$ is set by the association, the provider only has two rational choices: either to set $\sigma = \hat{\sigma}$ (but not higher) and seek certification, or to set $\sigma = 0$, forego certification, and save on the cost of producing data security.

This insight is used by the association when determining $\hat{\sigma}^*$. The association board makes a guess how a given level of $\hat{\sigma}$ would impact the incentives of the provider (i) to offer cloud services in the market altogether and (ii) to prefer seeking certification over non-certification. Importantly, although *ceteris paribus* users prefer a higher over a lower accountability level, because of the convex shape of the cost function to produce accountability, $c(\sigma)$, it is possible that producing highest accountability levels is so expensive for the provider that users would not be willing to pay the corresponding high price necessary to cover the provider’s cost. Therefore, if either $c(\sigma)$ or any one part determining the membership fee ($\frac{c_A}{\delta}$, c_L , and c_V) is too high as compared to the gross utility from consuming a secure cloud service (u), the mechanism breaks down and it is impossible to establish $\sigma > 0$ in a certification equilibrium.³⁶ Therefore, for the rest of this analysis, consider the case where those costs are not prohibitive.

A key characteristic of the cloud computing market is that increasing the credible accountability level of certified providers, $\hat{\sigma}$, drives up demand so much that the market is covered for all levels above some threshold (namely for $\hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}$). Equation (5) displays the provider’s profit-maximizing pricing strategy, which depends on this threshold. As long as demand is still elastic, the price is independent of $\hat{\sigma}$: If the association increases $\hat{\sigma}$ by one marginal unit, this

³⁵See Shapiro and Stiglitz (1984) for intuition and background on the efficiency wage hypothesis.

³⁶In this case, the model predicts that the cloud computing market is served by providers with very low accountability levels ($\sigma = 0$) meeting a positive, yet small amount of users paying a modest price— $p^*(x = 0) = \frac{u}{2} \Rightarrow \bar{s}_0 = \frac{u}{2} > 0$; see equations (A.9) and (A.5).

drives up demand for cloud services and thereby increases the provider’s revenues. By contrast, if $\hat{\sigma}$ reached the threshold level and all users buy cloud services, the provider’s best response to an additional increase in $\hat{\sigma}$ is to increase the price p^* , too.

This bifid equilibrium pricing strategy induces that the provider’s profit function has two local peaks: one in the range where demand is still elastic (characterized by a $\hat{\sigma}$ -level where the marginal cost of accountability, $c'(\hat{\sigma}) < \bar{s}$) and one in the range where the market is covered (characterized by a $\hat{\sigma}$ -level where $c'(\hat{\sigma}) = \bar{s}$). Which of these two local maxima is the global profit-maximizing $\hat{\sigma}$ -level depends on the shape of $c(\sigma)$ and on \bar{s} .³⁷

Finally, a crucial condition to establish the certification equilibrium characterized above is the requirement that the cost, c_C , of a user to complain to the cloud association in case she received cloud services with less accountability than certified is rather low. In this baseline model, Assumption 2 does the trick for us, assuming $c_C \rightarrow 0$. In Section 4, I discuss the consequences of significant positive costs of complaining. Independent of the level of c_C , however, the dependence of the certification equilibrium on c_C underlines one interesting feature of economic governance mechanisms. In order to obtain aspired behavior of agents, it is necessary to have credible punishment mechanisms in place, which are triggered only by unaspired behavior (here: certifiers and certified providers are punished only, by withdrawing business from them in the future, if they claim to produce accountability $\hat{\sigma}$ but then fail to do so). As soon as such trigger punishment is credible, however, it is not in the interest of agents to pull the trigger. Punishment is credible if it is in the interest of and affordable to the punisher. As a consequence, the actual punishment costs are saved. In this model, the costs that are saved on the equilibrium path are the cost of complaining (c_C) and the cost of ex post verification of the auditing procedure (c_V). The cost that are actually incurred by using the certification governance scheme are the cost of licensing certifiers (c_L) and the cost of auditing providers (c_A).

Further insights are generated by studying the equilibrium payoffs of provider and users, as detailed in Lemma A.5. As a group, users have a uniquely preferred accountability level. This level—in the model at $\hat{\sigma} = 1 - \frac{u}{2\bar{s}}$ —just leads to full market coverage: Already accounting for the provider’s pricing strategy, it is sufficient to convince the user with highest accountability preferences, at $s_i = \bar{s}$, to buy from the provider and obtain net indirect utility of zero. All other users receive positive consumer surplus (owing to the assumption that price discrimination according to individual accountability preferences is impossible).³⁸ The fact that users are

³⁷At the end of Appendix A, I provide a numerical example that shows how the model can be used to identify the $\hat{\sigma}^*$ -level producing the global profit-maximum.

³⁸If the association would increase $\hat{\sigma}$ even further, the net payoffs of users would decrease. This holds because the provider would increase the price in line with $\hat{\sigma}$, always making sure that the user at $s_i = \bar{s}$ is just indifferent between buying and not buying. Because all other users have a lower willingness-to-pay for marginal accountability, they would be charged the same marginal price increase targeted at the user at $s_i = \bar{s}$ but would receive lower marginal consumption utility because their $s_i < \bar{s}$. This process comes to an end, at the latest, where

represented on the cloud association board and equipped with the right to veto the strategy proposals of providers leads to the outcome that all users (but the indifferent one at $s_i = \bar{s}$) receive higher net utility if a credible certification scheme is in place than if it is not.

At equilibrium, users' payoffs only depend on their accountability preferences, proxied by \bar{s} , and on the threshold accountability level of certified providers, $\hat{\sigma}^*$, which depends on the cost of producing accountability. It is noteworthy that all costs of the certification process are borne by the provider.³⁹ This characteristic survives when there are two competing providers (see section 4.3). Consequently, it will be providers who stop using the certification scheme first if the usage costs grow too large.

4 Extensions and Robustness

4.1 The value of personalization technologies

Both Culnan and Armstrong (1999) and Chellappa and Sin (2005) find in survey data that consumers are willing to provide more information about themselves—and reap higher benefits from personalization—if they trust a vendor more to apply fair privacy protection procedures. By assuming u to be a constant in equation (1), I abstract from such gross utility gains through accepting lower privacy standards, for two reasons. First, independent of the amount of secret business data or personal data that becomes known to a cloud service provider via personalization technologies, the user will only be damaged if this data leaks to third parties—which is captured by the probability $(1 - \sigma)$ here. Second, I assume that the *net* utility of users (v_i) is decreasing in the extent of privacy infringements. If that were not true at least for some users, there would be no market for data protection, which is in contradiction to the evidence presented by Chellappa and Sin (2005), Solove (2007), Tsai et al. (2011), Goldfarb and Tucker (2012), and Pearson et al. (2012).

4.2 Certifying the cloud value chain

A key characteristic of the cloud computing industry, which contributes to the flexible supply of service providers and thereby to the potential efficiency gains of the entire industry, is that cloud service providers frequently subcontract parts of the service they offer customers to other service providers, who, in turn, may subcontract parts of the work even further, and so on (Haeberlen, 2010). Thereby, the provider modeled in section 3 above turns into a value chain of providers, which are legally connected by contracts. In this case, the spirit of the model would dictate that the original provider (coined *data controller*), who concludes a service level agreement with the user (coined *data owner*, in case of individual users potentially also *data*

$\hat{\sigma} = 1 - \frac{u^2}{4\bar{s}^2}$ is reached because there users are better off without certification.

³⁹This directly refers to $\frac{c_A}{\delta} + c_L + c_V$ but can also include c_C ; see Section 4.

subject), would be responsible for the accountability level of all upstream providers (coined *data processors*).⁴⁰ The data controller could fulfil this obligation by subcontracting only to other providers that are also certified by the cloud association and take responsibility that the data owner’s data never leaves the network of certified service providers.

Formally, if such a value chain consists of K service providers and σ_k is the accountability level of provider $k \in \{1, \dots, K\}$, at equilibrium we must have $\operatorname{argmin}\{\sigma_k \in \{\sigma_1, \dots, \sigma_K\}\} = \hat{\sigma}^*$; see equation (4). The weakest link of the chain must be sufficiently strong to sustain credibility in the entire chain.

4.3 The case of competition

Consider first the case of duopoly competition. If two cloud service providers play the game described above simultaneously, it can be shown that at equilibrium one provider will join the cloud association, pay the fee F , seek certification, produce high accountability $\sigma = \hat{\sigma}$, and command a high price in the market. The other provider will not join the association, produce no accountability ($\sigma = 0$) and charge a low but strictly positive price.⁴¹ Consequently, just as in Proposition 1, users with high accountability preferences $s_i > \bar{s}_1^*$ do not buy. In the duopoly, however, cloud services are vertically differentiated and, hence, the market is split: users with intermediate preferences, $s_i \in (\bar{s}_0^*, \bar{s}_1^*]$, buy from the certified provider, whereas users with low valuation for accountability, $s_i \leq \bar{s}_0^*$, buy from the uncertified provider; where $\bar{s}_0^* < \bar{s}_1^*$.⁴²

These results create the prediction that competition among cloud service providers, when credible certification is available, will lead to provider stratification. Users who do not value privacy a lot can get cheap cloud services but face a large risk that third parties access their data. Users with higher security demand can be relatively sure that their data is secure but have to pay a premium for such security.⁴³ One important requirement to obtain a certified provider in such a duopoly model is that the threshold accountability level to become certified, $\hat{\sigma}^*$, is not too low. Otherwise, the additional security that a certified provider can credibly supply over the uncertified outside option ($\hat{\sigma}^* - 0$) is too low to command a price that covers the costs of producing $\hat{\sigma}^*$ and paying the certification process.

If the number of competing cloud service providers is increased above two, there are still no more than two equilibrium accountability levels, $\hat{\sigma}^*$ and 0. Standard economic theory would predict that entry occurs as long as the equilibrium profits of all providers are zero. In the model set-up of this paper, given that there are no costs for uncertified providers, many firms

⁴⁰Pearson (2011) and Pearson et al. (2012) explain the terminology used in more detail.

⁴¹Proofs for these results are available from the author upon request.

⁴²See Shaked and Sutton (1982) and the subsequent literature for more findings on vertically differentiated products.

⁴³Users can never trust privacy in the cloud absolutely as long the residual risk of a security incident $(1 - \hat{\sigma}^*) > 0$. See section 4.1 and footnotes 28 and 31 related to this point.

could offer uncertified services and de facto find themselves in a market with Bertrand price competition with homogenous goods, driving the price for uncertified services (and $\sigma = 0$) to marginal cost, i.e. zero. This result does not extend to the market segment of certified cloud services. If more than one provider gets certification, total demand for that segment ($\bar{s}_1^* - \bar{s}_0^*$) is split among those providers, which reduces revenues of any individual provider. Such entry could be incentive-compatible for a limited number of certified providers at equilibrium, given the right parameter values. But as soon as demand is split too much due to entry of further certified providers, the certification equilibrium would break down because the fixed costs of getting certification could not be borne by associated revenues, anymore.

One way to delay market breakdown would be the introduction of several vertically stratified certification thresholds, each handing out different certificates, for example, at the *gold*, *silver*, and *bronze* levels (where $\hat{\sigma}_{gold} > \hat{\sigma}_{silver} > \hat{\sigma}_{bronze}$). Each level could sustain one (or a few) providers, depending on parameter realizations, and cut the entire market in several niches. The problem that such differentiation would create is, however, that it requires users to know about and understand all available niches, such that they can make informed consumption decisions. This requirement stands in contrast to one of the motivations of this paper, that many users have a limited understanding of the technical subtleties of cloud services, which renders the mechanism with one unique certification threshold studied in the baseline model very attractive.

4.4 Multiple certifiers and multiple providers

The baseline model studies one provider and one certifier. If there were N licensed certifiers who would be randomly matched with M providers demanding certification every period, the probability of each licensed certifier to get a job in a certain future period would turn to M/N . Consequently, the net present value from investing auditing effort $a = 1$ and honestly awarding certificates would turn to $\frac{1}{1-\delta} \frac{M}{N} (R - c_A)$. The threshold revenue to make this strategy incentive compatible for all N licensed certifiers become $R \geq \frac{(1-\delta)N + \delta M}{\delta M} c_A$, which is larger than the threshold revenue depicted in equation (A.7) of the baseline model if, and only if, $N > M$. Apart from this quantitative adjustment, the quality of the remaining results is robust up to the limits discussed in section 4.3.

4.5 Positive costs of complaining

Consider the case where Assumption 2 does not hold, that is, where a user who complains about low accountability of a provider, bears a substantial cost $c_C > 0$. In this case, Lemma A.1.(ii) states that only security-sensitive users (with $s_i \geq \frac{c_C}{\delta}$) would complain in equilibrium. This implies that if a provider produces $\sigma < \hat{\sigma}$, from all \bar{s}_1^* buyers a mass $\frac{c_C}{\delta}$ would not

complain. If c_C was sufficiently large, one might guess that this would invite some providers to gamble and produce low accountability and their certifiers to risk handing out certificates without investing any auditing effort. The incentive compatibility constraint of a certifier to invest $a = 1$ and award the certificate truthfully, equation (A.6), would become harder to fulfil because in expectation the certifier would get payoff R in every future period with probability $\frac{c_C}{\sigma}$ despite having cheated. This effectively increases the necessary efficiency wage R^* above the level given in equation (A.7).

If we allow for a slight adjustment of the mechanism, this ad hoc intuition is misleading, though. Alternatively to increasing the income of certifiers, R^* , the cloud association could just determine that a user who correctly complains about the low accountability level of a provider does not only get substitute services with accountability $\hat{\sigma}^*$ but also a lump-sum payment c_C . In order to have funds to pay out such lump-sum payments, the association would have to slightly increase the membership-fee F . Then all users would be incentivized to complain whenever they find $\sigma < \hat{\sigma}^*$. Because, on the equilibrium path, there is no fraud and, hence, no complaints, this promise would even come for free after all (see the explanation at the end of section 3).

5 Discussion, Implications, and Further Research

The huge upside potential that cloud computing technologies offer both to producers and users, teamed with significant impediments to realizing this potential because of users' lack of trust in the security of sensitive data they put to the cloud, got this study started. Such lack of trust is a consequence of several interrelated problems stemming from asymmetric information between sellers, buyers, and third parties supporting their transaction. Once a cloud service provider implemented a high accountability level, she suffers from an adverse selection problem because she has no means to credibly convince users that their data are secured. As soon as certification agencies, who audit providers' accountability levels and award "trust" certificates, are employed, such adverse selection can be mitigated but an additional problem arises: certifiers' moral hazard to actually spend unverifiable effort on understanding the true accountability levels of providers.

Inspired by the structure of existing certification schemes, I develop and apply the idea that representatives of both providers and users should jointly oversee the certification process in an independent, private body, called cloud association. Due to capacity constraints at the board level, this association sources the actual auditing process and certification decisions out to independent certification agencies. But it requires these agencies to undergo a costly licensing procedure ex ante, which enables it to verify ex post whether the auditor committed fraud. In turn, opportunistic behavior by the cloud association, in its function as auditor of auditors, is avoided because the governance structure of the association establishes checks and balances between representatives of both sides of the market.

My results show that in a dynamic, globally operating industry a private ordering institution, in the sense of Williamson (2002), can avoid market breakdown (in theory) and thereby support market growth (in practice). In Williamson’s (1991) terms, the problem of adaptation (to the dynamic, global market) can best be solved by a hybrid governance structure that combines hierarchical elements (within the association) and market elements (between the association and certifiers) without relying on public enforcement. The effectiveness of this institution depends on the careful composition of organizational and institutional features. For instance, in this model the profit-motive of both cloud service providers and certification agencies motivates them to act in the way aspired by the cloud association, whereas it is critical to take away the profit-motive from decision-makers within the association by incorporating it as a nonprofit organization and implementing checks and balances between the two camps of representatives of both market sides, providers and users.

Because the cloud association keeps records about the reported history of all certifiers, it serves as an institutionalized, central information repository of the cloud computing industry. This characteristic enables the association to play a grim trigger (ostracism) strategy with respect to certifiers: it can condition the award of a license to a certifier on her history and refuse to license a certifier who has ever falsely awarded a certificate, thereby expelling the fraudulent certifier from the community of licensed certifiers—and from all related future profits. This threat is credible because cheated users are incentivized to report their damage to the cloud association (via expecting indemnification for their losses from low accountability by receiving a substitute, high-accountability service) and because the requirement of using the association’s auditing protocol makes the actions of the certifier verifiable *ex post*. The “eternal memory” characteristic of the cloud association is a key advantage over any mechanism that only involves decentralized players, who suffer from imperfect transmission of information about a certifier’s past actions.

As proposed in this paper, the idea of the cloud governance scheme surrounding the cloud association is novel. But related schemes do exist in practice. For example, the Cloud Security Alliance (CSA) describes itself as “a not-for-profit organization with a mission to promote the use of best practices for providing security assurance within Cloud Computing, and to provide education on the uses of Cloud Computing to help secure all other forms of computing. The Cloud Security Alliance is led by a broad coalition of industry practitioners, corporations, associations and other key stakeholders” (<https://cloudsecurityalliance.org/about/>). Membership for corporate members, usually cloud service providers and corporate cloud customers, is available for U\$ 10,000 p.a. (as of November 2014) and gives the members voice within the CSA. A key task of the CSA is to define provider certification schemes (called Open Certification Framework (OCF)/STAR Registry) and auditor licensing schemes—analogous to setting $\hat{\sigma}$ spending c_L in my model. Under the STAR Certification rules, only nationally licensed certification bodies

can become STAR Certification Auditors—provided they become CSA Corporate Members, obtain special training, and pass an exam (resembling the licensing procedure and cost c_L in my model). The CSA decides when to review the certification bodies, a process that is paid for by the certifiers. The CSA is governed by a board, which hires an executive team of industry experts as managers.

The key differences between the existing CSA scheme and the scheme proposed in this paper are the following: First, the proposed scheme does not rely on governmental pre-licensing of certifiers by their national governments (because of the problems of public ordering discussed in the Section 2). Second, this scheme advocates a strong role of users, including decision-making power (or at least veto power) on the cloud association board, whereas the current CSA scheme is more open, allowing both corporate and individual members to raise their voices regarding key decisions; this cannot avoid, however, that user interests can be overridden by provider interests. Third, the CSA has still underdeveloped tools to receive and incorporate end user feedback and complaints about specific providers, let alone to incentivize users to file complaints if they think they were cheated.

Regarding the practical implications suggested by the model in this paper, it is critical to bear the conditions in Proposition 1 in mind under which the certification equilibrium can exist: the cost parameters ($c_C, c_A, c_L, c_V, c(\sigma)$) must not be too high and consumption utility u must not be too low. This predicts that only cloud services that create sufficient value for users will be accompanied by certification schemes. For low value applications, certification is too expensive. More importantly, users' cost of complaining, certifiers' cost of auditing according to the association's protocol, the association's cost of licensing and educating certifiers and verifying an actual certification procedure *ex post*, as well as provider's cost of producing accountability must not be too high if the certification scheme suggested here is to work in practice. This insight calls for the least complex certification procedure that can assure a certain accountability level, and it calls for the development of software systems that reduce all these costs, for instance to make complaining about a provider as simple and cheap as possible for users. The model also strongly suggests that public decision makers do not get directly involved in the procedures and enforcement of the cloud association's rules.

Although this model is designed for the cloud computing industry, it can inform trust-based governance schemes in other industries that are subject to asymmetric information problems as well. Its main power stems from the reduction of complexity from consumers' perspectives, which is achieved by the transformation of a multidimensional quality vector into an easily observable, binary certified/not certified-variable. This characteristic is shared by many high-tech industries with both business-to-business and business-to-consumer transactions. A further characteristic that facilitates the application of the cloud association governance structure to other industries is a similar cost structure, especially constant and low marginal costs of pro-

duction and distribution, such that capacity constraints are usually not binding and such that sellers sell their products worldwide to heterogeneous consumers. Nevertheless, it remains important to study the specific institutional background of the industry at stake in detail and not to plainly transplant the governance structure proposed here for the cloud. For instance, Montiel, Husted, and Christmann (2012) showed, based on ISO 14001 environmental management system certification among automotive plants in Mexico, that widespread corruption in the general environment can extend distrust to private certification systems, thereby reducing the credibility of private certification schemes. In such an environment, certification schemes may not be the optimal contract enforcement institution.

In this paper I have shown that and why a two-tiered certification framework designed around the cloud association can improve the use of cloud computing technologies. However, it is subject of future research both in management science and in information systems and computer science to study its optimal implementation. For inspiration, Albuquerque, Bronnenberg, and Corbett (2007) have compared diffusion patterns of ISO 9000 and ISO 14000 certification standards and shown which impact factors critically facilitate adoption there. Their results may be relevant for cloud computing, too. Insights from legal scholarship may be helpful in identifying the necessary legal circumstances when establishing the cloud association framework, and further research on the dynamics of organizational development are needed to better understand the way from where we stand today to a fully implemented governance institution. The cloud computing industry draws on many fields of expertise. As researchers, we may also need to cross disciplinary boundaries in order to help solve its problems and realize its full growth potential.

Appendix

A Proof of Proposition 1

Preliminary remarks: The six-stage game in period t described in the main text is solved by backward induction. In this appendix, I formally solve the model. Interpretation and economic intuition of the intermediate results is provided in the main text. Before detailing the solution, let me clarify the incentives of the provider and the resulting expectations of the users.

Call the accountability level announced by the provider σ' . The accountability level expected by a user is denoted $\tilde{\sigma}$. If no certification procedure is in place, the choice of σ made by the provider is unobservable to all other players. Because $c(\sigma)$ is increasing in σ , it is rational for the provider to set $\sigma = 0$ then. Users realize this incentive. Consequently, independent of σ' , without credible certification we have $\tilde{\sigma} = 0$ at equilibrium.

In turn, if a credible certification mechanism is in place (which is shown below) that hands out a certificate to a provider if, and only if, the provider's $\sigma \geq \hat{\sigma}$, excess accountability above the threshold level is still unobservable and hence would not be rewarded by paying a higher price by users. It follows that, if a provider wants to earn a certificate and expects the auditor to make truthful decisions, her optimal response is to set $\sigma = \hat{\sigma}$. Denote the fact that a provider is certified by $x = 1$ and the fact that a provider is uncertified by $x = 0$. Then, if a credible certification mechanism is in place, users' accountability beliefs are:

$$\tilde{\sigma} = \begin{cases} \hat{\sigma} & \text{if } x = 1, \\ 0 & \text{if } x = 0. \end{cases} \quad (\text{A.1})$$

Proof: At stage 6, if a user received a service with accountability level $\sigma \geq \hat{\sigma}$, a complaint makes no sense because complaining is costly and ex post verification of the facts by the cloud association is perfect. Thus, the expected payoff from complaining in this case would be $-c_C$. In contrast, if a user received a service with accountability level $\sigma < \hat{\sigma}$, the user expects to be reimbursed by the cloud association with a substitute service of accountability $\hat{\sigma}$. Consequently, the expected payoff from filing a complaint and receiving the substitute service is $u - (1 - \hat{\sigma})s_i - c_C$. By contrast, without filing a complaint, the expected payoff remains $u - (1 - \sigma)s_i$. It is in the user's interest to file a complaint if, and only if:

$$u - (1 - \hat{\sigma})s_i - c_C \geq u - (1 - \sigma)s_i \quad (\text{A.2})$$

$$\Leftrightarrow c_C \leq (\hat{\sigma} - \sigma)s_i \quad (\text{A.3})$$

Because, by (A.1), the unique equilibrium accountability without certification is $\sigma = 0$, we find:

Lemma A.1 (Equilibrium user complaints) (i) *If, and only if, the cost of complaining is sufficiently small, cheated users will complain to the association (if $c_C \leq \hat{\sigma}s_i \equiv \bar{c}_C$).* (ii) *Only sufficiently security-sensitive users will complain (those for whom $s_i \geq \frac{c_C}{\hat{\sigma}}$).*

Because of Assumption 2, these constraints are fulfilled for approximately all users.⁴⁴

At stage 5, every user can decide whether to buy from the monopolistic provider, or not. Buying gives expected indirect utility of $u - (1 - \tilde{\sigma})s_i - p$. Because of (A.1), it follows that all users with accountability preferences

$$s_i \leq \frac{u - p_1}{1 - \hat{\sigma}} \equiv \bar{s}_1 \quad (\text{A.4})$$

buy cloud services if $x = 1$, whereas all users with preferences

$$s_i \leq u - p_0 \equiv \bar{s}_0 \quad (\text{A.5})$$

buy cloud services if $x = 0$, where p_1 (p_0) denotes the price charged by a certified (uncertified) provider and $\bar{s}_1 > \bar{s}_0$. As the maximum of \bar{s}_1 is \bar{s} , the market is fully covered if $p_1 \leq u - (1 - \hat{\sigma})\bar{s} \Leftrightarrow \hat{\sigma} \geq 1 - \frac{u - p_1}{\bar{s}}$, or if $p_0 \leq u - \bar{s}$.

At stage 4, the commissioned certification agency decides about auditing effort a and certification x . Because of the cloud association's ability to use a grim trigger strategy, that is, to exclude shirking certifiers from all business in the future, the certifier's trade-off is as follows: If it invests the auditing cost c_A and truthfully awards the certificate, it expects a net payoff, $R - c_A$, now and in every subsequent period t in which it is commissioned by the cloud association to conduct an audit. In this baseline model with 1 certifier, it expects an audit job every future period. If the certifier does not invest in auditing, it can only make a guess about the provider's true σ . Recall that the solution concept used is subgame-perfect equilibrium in pure strategies. This implies that, if the certifier would not spend c_A but still award a certificate, the provider's best-response would be to set $\sigma = 0$. In this case, however, because of Lemma A.1, the fraud would be detected and the certifier would never obtain an auditing job in the future again.⁴⁵ Consequently, the certifier will invest auditing effort and truthfully decide about x if, and only if:

$$\frac{1}{1 - \delta}(R - c_A) \geq R \quad (\text{A.6})$$

$$\Leftrightarrow R \geq \frac{c_A}{\delta} \quad (\text{A.7})$$

The left-hand side (LHS) of (A.6) is the certifier's net present value if she invests effort into auditing and truthfully awards the certificate; the right-hand side (RHS) is the net present value from cheating once, saving c_A , and never earning R in the future. Re-arranging yields (A.7).

At stage 3, the cloud association commissions the licensed certifier and determines the revenue R it offers for conducting the audit.⁴⁶ Because the association is run on behalf of

⁴⁴See Section 4.5 for the adapted results and a simple solution if complaining is costly.

⁴⁵If the certifier sets $a = 0$ and does not award a certificate, she risks that a provider with $\sigma = \hat{\sigma}$ complains to the association, which, after verification, implies a loss of all future auditing mandates for the certifier.

⁴⁶See section 4.4 for the case where one certifier is randomly chosen from a pool of licensed certifiers.

providers and users and not on behalf of certifiers, the nonprofit association's board will minimize R conditional on the commissioned certifier accepting the task and being incentivized to act truthfully. Hence, we obtain the following Lemma.

Lemma A.2 (Equilibrium certifier revenues and auditing decisions) *(i) The cloud association offers a certifier revenue $R^* \equiv \frac{c_A}{\delta}$ per audit. (ii) The certifier accepts the auditing job, invests effort $a = 1$, and decides truthfully about the provider's certification status.*

At stage 2, the provider sets accountability level σ and price p , and decides whether to spend F and seek certification by becoming an association member. Following (A.4) and (A.5), the provider faces demand $q_1 = \bar{s}_1/\bar{s}$ if $x = 1$ and $q_0 = \bar{s}_0/\bar{s}$ if $x = 0$. Given (A.1), seeking certification implies setting $\sigma = \hat{\sigma}$ and not seeking certification implies $\sigma = 0$. Because we have to distinguish among cases where the market is covered and where demand is elastic and cases with and without certification, the provider's profit-maximization problem, (2), can be rewritten as:

$$\max \pi = \begin{cases} p \frac{u-p}{(1-\hat{\sigma})\bar{s}} - c(\hat{\sigma}) - F & \text{if } \sigma = \hat{\sigma} \text{ and } p > u - (1 - \hat{\sigma})\bar{s}, \\ p - c(\hat{\sigma}) - F & \text{if } \sigma = \hat{\sigma} \text{ and } p \leq u - (1 - \hat{\sigma})\bar{s}, \\ p \frac{u-p}{\bar{s}} & \text{if } \sigma = 0 \text{ and } p > u - \bar{s}, \\ p & \text{if } \sigma = 0 \text{ and } p \leq u - \bar{s}. \end{cases} \quad (\text{A.8})$$

The resulting equilibrium price, depending on certification and parameter realizations, is:⁴⁷

$$p^* = \begin{cases} \frac{u}{2} & \text{if } \sigma = \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \\ u - (1 - \hat{\sigma})\bar{s} & \text{if } \sigma = \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}, \\ \frac{u}{2} & \text{if } \sigma = 0 \text{ and } \frac{u}{2} < \bar{s}, \\ u - \bar{s} & \text{if } \sigma = 0 \text{ and } \frac{u}{2} \geq \bar{s}. \end{cases} \quad (\text{A.9})$$

Assumption 1 rules out the fourth row of equation (A.9). Substituting (A.9) in (A.8) yields the provider's equilibrium profits.

$$\pi^* = \begin{cases} \frac{u^2}{4(1-\hat{\sigma})\bar{s}} - c(\sigma) - F & \text{if } \sigma = \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \\ u - (1 - \hat{\sigma})\bar{s} - c(\sigma) - F & \text{if } \sigma = \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}, \\ \frac{u^2}{4\bar{s}} & \text{if } \sigma = 0 \end{cases} \quad (\text{A.10})$$

⁴⁷This result is qualitatively robust to introducing a constant marginal cost of production, $k > 0$, which increases the equilibrium price if demand is flexible to $P = \frac{u+k}{2}$. It adapts if the market structure changes. See section 4.3 for the case of duopolistic providers.

The provider prefers to seek certification and set accountability $\sigma = \hat{\sigma}$ over not seeking certification and setting $\sigma = 0$ if, and only if:

$$\frac{u^2}{4(1-\hat{\sigma})\bar{s}} - c(\sigma) - F \geq \frac{u^2}{4\bar{s}} \quad \text{if } \sigma = \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \quad (\text{A.11})$$

$$u - (1-\hat{\sigma})\bar{s} - c(\sigma) - F \geq \frac{u^2}{4\bar{s}} \quad \text{if } \sigma = \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}. \quad (\text{A.12})$$

Rearranging (A.11) and (A.12), certification is attractive for the provider if, and only if:

$$\hat{\sigma} \geq 1 - \frac{u^2}{4\bar{s}(c(\sigma) + F) + u^2} \quad \text{if } \sigma = \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \quad (\text{A.13})$$

$$\hat{\sigma} \geq \frac{2(\bar{s}-u)^2}{4\bar{s}^2} + \frac{c(\sigma) + F}{\bar{s}} \quad \text{if } \sigma = \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}. \quad (\text{A.14})$$

Equation (A.13) determines a non-empty set if $1 - \frac{u^2}{4\bar{s}(c(\sigma)+F)+u^2} < 1 - \frac{u}{2\bar{s}}$. Hence, seeking certification can be made attractive for the provider, via determining $\hat{\sigma}$, if and only if:

$$c(\sigma) + F < \frac{u}{2} + \frac{u^2}{4\bar{s}} \quad \text{for } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \quad (\text{A.15})$$

In turn, both conditions in equation (A.14) describe lower bounds on $\hat{\sigma}$. The binding (higher) condition for the case where $\hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}$ is:

$$\hat{\sigma} \geq \begin{cases} 1 - \frac{u}{2\bar{s}} & \text{if } u \in (0, \frac{1}{2}(\bar{s} - \sqrt{5\bar{s}^2 - 8\bar{s}(c(\sigma) + F)})) \\ \frac{2(\bar{s}-u)^2}{4\bar{s}^2} + \frac{c(\sigma)+F}{\bar{s}} & \text{if } u \in [\frac{1}{2}(\bar{s} - \sqrt{5\bar{s}^2 - 8\bar{s}(c(\sigma) + F)}, 2\bar{s}) \end{cases} \quad \text{and } c(\sigma) + F < \frac{5}{8}\bar{s} \quad (\text{A.16})$$

Lemma A.3 (Provider incentives and pricing) *The provider prices cloud services as in (A.9). She sets $\sigma = \hat{\sigma}$, joins the association for the fee F , and seeks certification if the respective conditions in (A.15) and (A.16) hold.*

At stage 1, the cloud association's board determines $\hat{\sigma}$ and F and licenses a certifier, for cost c_L . As F is the association's only source of income, it must be able to cover all its expenses related to the licensing of one certifier and the certification of one provider, namely R^* , c_L , and c_V .⁴⁸ Using Lemma A.2, the lower bound on F (the minimum membership-fee) is $\underline{F} = \frac{c_A}{\delta} + c_L + c_V$. Neither the representatives of providers nor those of users on the association board have an interest in charging excessive fees, which would only make association membership less attractive but serve no positive purpose. Hence \underline{F} equals the equilibrium fee.

Regarding $\hat{\sigma}$, the preferences of the board representatives of providers and users are different, however. Increasing $\hat{\sigma}$ beyond $(1 - \frac{u}{2\bar{s}})$ decreases the provider's profits because revenues stay constant but accountability costs increase. It follows that the representatives of providers solve

⁴⁸To demonstrate the main mechanism clearly, I abstract away both from factors that could influence F in case of multiple certifiers or providers, such as economies of scale in licensing or the fact that on the equilibrium path spending c_V will be unnecessary.

the following program:

$$\begin{aligned} \max_{\hat{\sigma}} \quad \pi = & \begin{cases} \frac{u^2}{4(1-\hat{\sigma})\bar{s}} - c(\hat{\sigma}) - \underline{F} & \text{if } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \\ u - (1-\hat{\sigma})\bar{s} - c(\sigma) - \underline{F} & \text{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}. \end{cases} \quad (\text{A.17}) \\ \text{s.t.} \quad & (\text{A.15}) \text{ and } (\text{A.16}) \text{ hold and } \pi(x=1) \geq 0. \end{aligned}$$

The side-constraints make sure that (i) seeking certification is incentive compatible for the provider and (ii) that the provider's participation constraint (to produce without making losses) holds. Because the provider makes positive profits without certification—see (A.10)—if the incentive compatibility constraints, (A.15) and (A.16), hold, the participation constraint also holds (and hence does not have to be considered explicitly any further).

To find the most-preferred (=profit-maximizing) level of $\hat{\sigma}$ for the provider we solve (A.17) and find that it has a local maximum at:

$$\hat{\sigma} = \begin{cases} 1 - \frac{u}{2\sqrt{\bar{s}c'(\sigma)}} & \text{if } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \\ \sigma | c'(\sigma) = \bar{s} & \text{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}. \end{cases} \quad (\text{A.18})$$

where $c'(\sigma)$ denotes the provider's marginal cost of accountability. The first row of (A.18) implies $c'(\sigma) < \bar{s}$, the second row implies $c'(\sigma) = \bar{s}$. It follows that, depending on the shape of $c(\sigma)$, provider representatives prefer to set $\hat{\sigma} = \{\sigma | c'(\sigma) < \bar{s}\}$ until the threshold level $\hat{\sigma} = 1 - \frac{u}{2\bar{s}}$ is reached. Then $\hat{\sigma} = \{\sigma | c'(\sigma) = \bar{s}\}$ is profit-maximizing. Which of the two $\hat{\sigma}$ -levels solves (A.17) depends on the shape of $c(\sigma)$. This completely characterizes the solution to the providers' representatives' optimization problem.

The second party on the cloud association board are user representatives, who maximize expected consumer surplus:

$$CS = \begin{cases} \int_0^{q_1} (u - (1-\hat{\sigma})s_i - p_1) ds_i & \text{if } x = 1, \\ \int_0^{q_0} (u - s_i - p_0) ds_i & \text{if } x = 0. \end{cases} \quad (\text{A.19})$$

Substituting $q_1 = \bar{s}_1/\bar{s}$, $q_0 = \bar{s}_0/\bar{s}$, and p^* from (A.9) we obtain:

$$CS = \begin{cases} \frac{(2\bar{s}-1)u^2}{8\bar{s}^2(1-\hat{\sigma})} & \text{if } \sigma = \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \\ \frac{1}{2}(2\bar{s}-1)(1-\hat{\sigma}) & \text{if } \sigma = \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}, \\ \frac{(2\bar{s}-1)u^2}{8\bar{s}^2} & \text{if } \sigma = 0. \end{cases} \quad (\text{A.20})$$

Users prefer certification and $\sigma = \hat{\sigma}$ over no certification and $\sigma = 0$ whenever:

$$\frac{(2\bar{s}-1)u^2}{8\bar{s}^2(1-\hat{\sigma})} > \frac{(2\bar{s}-1)u^2}{8\bar{s}^2}, \quad \text{if } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \quad (\text{A.21})$$

$$\frac{1}{2}(2\bar{s}-1)(1-\hat{\sigma}) > \frac{(2\bar{s}-1)u^2}{8\bar{s}^2}, \quad \text{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}} \quad (\text{A.22})$$

Inequality (A.21) holds for all $\hat{\sigma} > 0$. The difference between the LHS and RHS of (A.21) is growing in $\hat{\sigma}$. Hence, in this case the user representatives prefer the highest supported level

of $\hat{\sigma}$, for which $\hat{\sigma} < 1 - \frac{u}{2\bar{s}}$. Inequality (A.22) only holds for $\hat{\sigma} \in [1 - \frac{u}{2\bar{s}}, 1 - \frac{u^2}{4\bar{s}^2}]$. In this range, consumer surplus from certification is monotonically decreasing in $\hat{\sigma}$. Hence, the user representatives prefer the lowest supported level of $\hat{\sigma}$, $1 - \frac{u}{2\bar{s}}$. Here is a summary of the stage 1 results.

Lemma A.4 (Equilibrium membership-fee and accountability preferences) (i) The equilibrium membership-fee is $F^* = \frac{c_A}{\delta} + c_L + c_V$. (ii) The board representatives of users would prefer an accountability level of $\hat{\sigma}_U^* = 1 - \frac{u}{2\bar{s}}$ but will not veto any $\hat{\sigma} \leq 1 - \frac{u^2}{4\bar{s}^2}$. (iii) The provider's representatives prefer the threshold accountability level to be:

$$\hat{\sigma}_P^* = \begin{cases} \hat{\sigma}_1 & \text{if } \frac{u^2}{4(1-\hat{\sigma}_1)\bar{s}} - c(\hat{\sigma}_1) > u - (1 - \hat{\sigma}_2)\bar{s} - c(\hat{\sigma}_2), \\ \hat{\sigma}_2 & \text{otherwise} \end{cases} \quad (\text{A.23})$$

where $\hat{\sigma}_1 \equiv \{1 - \frac{u}{2\sqrt{\bar{s}c'(\sigma)}} | \sigma < 1 - \frac{u}{2\bar{s}}\}$ and $\hat{\sigma}_2 \equiv \{\sigma | c'(\sigma) = \bar{s} \wedge \sigma \geq 1 - \frac{u}{2\bar{s}}\}$. (iv) The equilibrium threshold accountability level of the cloud association is $\hat{\sigma}^* = \hat{\sigma}_P^*$ s.t. $\hat{\sigma}^* \leq 1 - \frac{u^2}{4\bar{s}^2}$.

Substituting $\hat{\sigma}^*$ and p^* in (A.4) we derive users' equilibrium consumption decisions (at stage 5) and the associated payoffs.

Lemma A.5 (Equilibrium consumption and payoffs) (i) In the certification equilibrium, all users with accountability preferences $s_i \leq \bar{s}_1^*$ buy the provider's service for the price p^* , where $\bar{s}_1^* = \frac{u}{2(1-\hat{\sigma}^*)}$ if $c'(\hat{\sigma}^*) < \bar{s}$, and $\bar{s}_1^* = \bar{s}$ if $c'(\hat{\sigma}^*) \geq \bar{s}$. (ii) Consumer surplus amounts to:

$$CS = \begin{cases} \frac{(2\bar{s}-1)u^2}{8\bar{s}^2(1-\hat{\sigma}^*)} & \text{if } c'(\hat{\sigma}^*) < \bar{s}, \\ \frac{1}{2}(2\bar{s}-1)(1-\hat{\sigma}^*) & \text{if } c'(\hat{\sigma}^*) \geq \bar{s}. \end{cases} \quad (\text{A.24})$$

(iii) The provider obtains surplus of:

$$\pi^* = \begin{cases} \frac{u^2}{4(1-\hat{\sigma}^*)\bar{s}} - c(\hat{\sigma}^*) - (\frac{c_A}{\delta} + c_L + c_V) & \text{if } c'(\hat{\sigma}^*) < \bar{s}, \\ u - (1 - \hat{\sigma}^*)\bar{s} - c(\hat{\sigma}^*) - (\frac{c_A}{\delta} + c_L + c_V) & \text{if } c'(\hat{\sigma}^*) \geq \bar{s}. \end{cases} \quad (\text{A.25})$$

(iv) Total welfare is given by:

$$W^* = \begin{cases} \frac{(4\bar{s}-1)u^2}{8\bar{s}^2(1-\hat{\sigma}^*)} - c(\hat{\sigma}^*) - (\frac{c_A}{\delta} + c_L + c_V) & \text{if } c'(\hat{\sigma}^*) < \bar{s}, \\ u - \frac{(1-\hat{\sigma}^*)}{2} - c(\hat{\sigma}^*) - (\frac{c_A}{\delta} + c_L + c_V) & \text{if } c'(\hat{\sigma}^*) \geq \bar{s}. \end{cases} \quad (\text{A.26})$$

Total welfare is equal where both cases merge, at $\hat{\sigma} = 1 - \frac{u}{2\bar{s}}$, namely there $W^* = u - \frac{u}{4\bar{s}} - c(\hat{\sigma}^*) - (\frac{c_A}{\delta} + c_L + c_V)$. This concludes the proof of Proposition 1. *Q.E.D.*

Illustrative example: For illustration, consider the following numerical example. Assume $u = \bar{s} = 1$, which satisfies Assumption 1. Assume $\frac{c_A}{\delta} + c_V + c_L = 0.1 = F$, and $c(\sigma) = \sigma^2$, which satisfies (3). Thus, we know that demand is elastic for all $\hat{\sigma} < 1 - \frac{u}{2\bar{s}} = \frac{1}{2}$. In this case,

certification is only attractive for the provider if (A.15) holds, which requires: $\hat{\sigma} + 0.1 < \frac{1}{2} + \frac{1^2}{4*1} \Leftrightarrow \hat{\sigma} < 0.65$. (A.18) then implies that the profit-maximizing $\hat{\sigma} = 1 - \frac{1}{2\sqrt{1*2\hat{\sigma}}} \Rightarrow \hat{\sigma}_{elastic}^* = 0.19$. Substituting values in (A.17) shows that $\pi_{elastic}^* = 0.17$. By contrast, demand is fixed at $q = 1$ for $\hat{\sigma} \geq \frac{1}{2}$. (A.18) states that the provider's profit is maximized at $\hat{\sigma}|c'(\hat{\sigma}) = \bar{s} \Leftrightarrow 2\hat{\sigma} = 1 \Rightarrow \hat{\sigma}_{fixed}^* = 0.5$. Substituting values in (A.17) shows that even with this profit-maximizing choice of $\hat{\sigma}$ and before constraining the choice further by applying (A.16), $\pi_{fixed}^* = 1 - (1 - \frac{1}{2}) * 1 - 2 * 0.5 - 0.1 = -0.6$. Finally, because $\pi_{elastic}^* > \max\{\pi_{fixed}^*, 0\}$, $\hat{\sigma}_P^* = \hat{\sigma}_{elastic}^*$ maximizes the provider's profits.

Consumers, on the other hand, prefer $\hat{\sigma}_U^* = 1 - \frac{u}{2\bar{s}} = 0.5$, which would generate $CS = \frac{1}{2}(2 * 1 - 1)(1 - 0.5) = 0.25$. By contrast, at $\hat{\sigma}_P^*$, $CS = \frac{(2*1-1)*1^2}{8*1^2(1-0.19)} = 0.15$.

References

- Akerlof, George. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84 (3): 488-500.
- Albuquerque, Paulo, Bronnenberg, Bart J. and Charles J. Corbett. 2007. "A Spatio-Temporal Analysis of the Global Diffusion of ISO 9000 and ISO 14000 Certification," *Management Science*, 53(3): 451-468.
- Aperjis, C and R. Johari. 2010. "Optimal windows for aggregating ratings in electronic marketplaces." *Management Science* 56(5): 864-880.
- Bae, Kee-Hong, Purda, Lynnette, Welker, Michael and Ligang Zhong. 2013. "Credit rating initiation and accounting quality for emerging-market firms," *Journal of International Business Studies*, 44: 216-234.
- Bakos, Yannis and Chrysanthos Dellarocas. 2011. "Cooperation Without Enforcement? A Comparative Analysis of Litigation and Online Reputation as Quality Assurance Mechanisms." *Management Science*, 57(11): 1944-1962.
- Buehler, Benno and Florian Schuett. 2014. "Certification and minimum quality standards when some consumers are uninformed." *European Economic Review*, 70: 493-511.
- Cabral, Luis M.B. and Ali Hortacsu. 2010. The dynamics of seller reputation: Theory and evidence from eBay. *Journal of Industrial Economics* 58(1): 54-78.
- Cai, Hongbin, Jin, Ginger Z., Liu, Chong and Li-An Zhou. 2013. "More Trusting, Less Trust? An Investigation of Early E-Commerce in China." NBER working paper 18961.
- Catteddu, Daniele and Giles Hogben. 2009. "An SME perspective on Cloud Computing," ENISA survey.
- Cezar, Asunur, Cavusoglu, Huseyin and Srinivasan Raghunathan. 2014. "Outsourcing Information Security: Contracting Issues and Security Implications." *Management Science* 60(3):638-657.
- Chellappa, Ramnath K. and Raymond G. Sin. 2005. "Personalization versus Privacy: An Empirical Examination of the Online Consumer's Dilemma." *Information Technology and Management* 6: 181-202.
- Christmann, Petra and Glen Taylor. 2006. "Firm Self-Regulation through International Certifiable Standards: Determinants of Symbolic versus Substantive Implementation." *Journal of International Business Studies*, 37(6): 863-878.
- Corbett, C.J., Montes-Sancho, M.J., and D.A. Kirsch. 2005. "The Financial Impact Of ISO 9000 Certification: An Empirical Analysis." *Management Science*, 51(7): 1046-1059.
- Culnan, M.J. and P.K. Armstrong. 1999. "Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation." *Organization Science* 10(1): 104-115.
- Dellarocas, Chrysanthos. 2003. "The digitization of word of mouth: Promise and challenges of online feedback mechanisms." *Management Science* 49(10): 1407-1424.
- Dellarocas, Chrysanthos and Charles A. Wood. 2008. "The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias." *Management Science* 54(3): 460-476.
- Dixit, Avinash. 2003. "Trade Expansion and Contract Enforcement," *Journal of Political Economy*, 111(6): 1293-1317.
- Dixit, Avinash. 2009. "Governance Institutions and Economic Activity," *American Economic Review*, 99(1): 5-24.
- Edelman, Benjamin. 2011. "Adverse selection in online 'trust' certifications and search results," *Electronic Commerce Research and Applications*, 10, 17-25.
- ENISA. 2013. "Schemes for Auditing Security Measures: An Overview," European Union Agency for Network and Information Security.
- Etro, Federico. 2009. "The Economic Impact of Cloud Computing on Business Creation, Employment and Output in Europe," *Review of Business and Economics*, 54(2): 179-208.

- EU Commission. 2014. "Internet Policy and Governance. Europe's role in shaping the future of Internet Governance." *Communication from the Commission to The European Parliament, The Council, The European Economic and Social Committee and The Committee of the Regions*. Brussels: COM(2014) 72/4.
- Filistrucchi, Lapo and Jens Prüfer. 2013. "Faithful Strategies: How Religion Shapes Nonprofit Management," Tilburg University, CentER Discussion Paper, No. 2013-052.
- Gibbons, Robert and Rebecca Henderson. 2012. "Relational Contracts and Organizational Capabilities," *Organization Science*, 23(5): 1350-1364.
- Goldfarb, Avi and Catherine Tucker. 2012. "Shifts in Privacy Concerns," *American Economic Review: Papers and Proceedings*, 102: 349-353.
- Greif, Avner, Milgrom, Paul R. and Barry R. Weingast. 1994. "Coordination, Commitment, and Enforcement: The Case of the Merchant Guild," *Journal of Political Economy*, 102: 745-776.
- Haeberlen, Andreas. 2010. "A Case for the Accountable Cloud," *ACM SIGOPS Operating Systems Review*, 44(2), 52-57.
- Hambrick, D.C. and P.A. Mason. 1984. "Upper Echelons: The Organization as a Reflection of Its Top Managers," *Academy of Management Review*, 9(2): 193-206.
- Li, Lingfang (Ivy) and Erte Xiao. 2014. "Money Talks: Rebate Mechanisms in Reputation System Design," *Management Science*, 60(8), 2054-2072.
- Masten, Scott E. and Jens Prüfer. 2011. "On the Evolution of Collective Enforcement Institutions: Communities and Courts," CentER Discussion Paper 2011-74.
- Masten, Scott E. and Jens Prüfer. 2014. "On the Evolution of Collective Enforcement Institutions: Communities and Courts," *Journal of Legal Studies*, 43(2), 359-400.
- Mathis, Jerome, McAndrews, James and Jean-Charles Rochet. 2009. "Rating the raters: Are reputation concerns powerful enough to discipline rating agencies?" *Journal of Monetary Economics*, 56: 657-674.
- Montiel, Ivan, Husted, Bryan W. and Petra Christmann. 2012. "Using Private Management Standard Certification to Reduce Information Asymmetries in Corrupt Environments." *Strategic Management Journal*, 33: 1103-1113.
- National Institute of Standards and Technology. 2009. "The NIST Definition of Cloud Computing," Information Technology Laboratory.
- New York Times. 2012. "New European Guidelines to Address Cloud Computing." July 1, 2012.
- O'Hara, E.A. 2005. "Choice of Law for Internet Transactions: The Uneasy Case for Online Consumer Protection." *University of Pennsylvania Law Review*, 153:1882-1950.
- Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge, UK: Cambridge University Press.
- Pearson, Siani. 2011. "Toward Accountability in the Cloud," *IEEE Internet Computing*, 15(4):64-69.
- Pearson, Siani, Catteddu, Daniele, Felici, Massimo, Hogben, Giles, Papanikolaou, Nick, and D. Pradelles. 2012. "Scoping report and initial glossary," A4Cloud project report, version 0.71.
- Peyrache, Eloic and Lucia Quesada. 2011. "Intermediaries, Credibility and Incentives to Collude," *Journal of Economics & Management Strategy*, 20(4): 1099-1133.
- Potoski, Matthew and Aseem Prakash. 2009. "Information Asymmetries as Trade Barriers: ISO 9000 Increases International Commerce," *Journal of Policy Analysis and Management*, 28(2): 221-238.
- Probst, Christian W., Sasse, M. Angela, Pieters, Wolter, Dimkov, Trajce, Luysterborg, Erik and Michel Arnaud. 2012. "Privacy Penetration Testing: How to Establish Trust in Your Cloud Provider," *European Data Protection: In Good Health?:* 251-265.
- Prüfer, Jens. 2013. "How to Govern the Cloud?" *IEEE CloudCom 2013*, DOI 10.1109/Cloud-Com.2013.100: 33-38.

- Prüfer, Jens. 2014. "Business Associations and Private Ordering," CentER Discussion Paper, No. 2012-094.
- Rao, Hayagreeva. 1994. "The Social Construction of Reputation: Certification Contests, Legitimation, and the Survival of Organizations in the American Automobile Industry: 1895-1912," *Strategic Management Journal*, 15: 29-44.
- Reed, Chris. 2013. "Governance in Cloud Computing," Queen Mary University of London, School of Law, Legal Studies Research Paper No. 157/2013.
- Shaked, Avner and John Sutton. 1982. "Relaxing Price Competition Through Product Differentiation," *Review of Economic Studies*, 49: 3-13.
- Shapiro, Carl, and Joseph E. Stiglitz. 1984. "Equilibrium Unemployment as a Worker Discipline Device," *American Economic Review*, 74(3): 433-444.
- Soghoian, Christopher. 2010. "Caught in the Cloud: Privacy, Encryption, and Government Back Doors in the Web 2.0 Era," *Journal on Telecommunications and High Technology Law*, 8(2).
- Solove, Daniel J. 2007. "'I've Got Nothing to Hide' and Other Misunderstandings of Privacy", *San Diego Law Review*, 44: 745-772.
- Spence, Michael. 1973. "Job Market Signaling," *Quarterly Journal of Economics*, 87(3): 355-374.
- Stahl, Konrad and Roland Strausz. 2011. "Who should pay for certification?," ZEW Discussion Papers, No. 11-054.
- Strausz, Roland. 2005. "Honest certification and the threat of capture." *International Journal of Industrial Organization*, 23: 45-62.
- Tsai, J., Egelman, S., Cranor, L. and A. Acquisti. 2011. "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study," *Information Systems Research*, 22: 254-268.
- Williamson, Oliver E. 1991. "Comparative Economic Organization: The Analysis of Discrete Structural Alternatives," *Administrative Science Quarterly*, 36(2): 269-296.
- Williamson, Oliver E. 2002. "The Lens of Contract: Private Ordering," *American Economic Review P&P*, 92(2): 438-443.