Trusting Privacy in the Cloud [*]

Jens Prüfer

Tilburg University[†]

October 29, 2018

Abstract

Cloud computing technologies can increase innovation and economic growth considerably. Because of privacy concerns, however, many users underutilize cloud technologies. This paper designs an institution attenuating the problem: a two-layered certification scheme built around a private, nonprofit organization called *cloud association*. This association is governed by representatives of both users and cloud service providers and sources auditing and certification of providers out to independent for-profit certifiers. It is shown how this institution incentivizes providers to produce high data security, and users with strong privacy preferences to trust them and pay a premium for their services.

JEL classification: D02, D71, L14, L31, L86

Keywords: trust; private ordering; associations; certification; privacy

[*] I am grateful to several seminar audiences at Tilburg University, at the A4Cloud General Meeting in Tilburg, the CPB in Den Haag, 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.

[†] Department of Economics, 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 2009). 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). Global cloud market revenues are predicted to increase from U$180b in 2015 to U$390b in 2020, attaining 17% annual average growth.[1] All indicators for the scope of the cloud computing industry, such as data traffic, the amount of data centers or the amount of cloud services are also predicted to grow exponentially (Cisco 2018).

But there are important impediments to the wide adoption of cloud computing. Most importantly, 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 (August et al. 2014).[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 cloud computing the key question is not about the optimal level of privacy regulation, that is, how much data on consumers' behavior should a provider be allowed to collect and use legally.[4] Instead, the key problem to be studied in this context is how to *enforce* cloud service

[1] `http://www.forbes.com/sites/louiscolumbus/2017/02/11/global-cloud-spending-predicted-to-reach-390b-by-2020/`

[2] A recent survey among IT decision makers in companies states that only 23% of organizations completely trust public clouds to keep their data secure (McAfee 2017). Among individual users, the hacking of iCloud pictures of celebrities is but one event that sparked fears for privacy (`http://www.computerworld.com/article/2601905/apple-icloud-take-reputation-hits-after-photo-scandal.html`).

[3] In Europe, sales of cloud computing services trail those in the US by at least two years, owing to concerns about privacy and fears that business secrets could be stolen in US based cloud centers (New York Times 2012). Data from a 2016 survey among cloud users underlines the topicality of the US-EU divide in planned cloud usage because of privacy concerns (`https://iapp.org/news/a/on-the-importance-of-data-privacy-to-cloud-deals-some-benchmarks/`).

[4] See Taylor (2004), Acquisti and Varian (2005), Calzolari and Pavan (2006), Casadesus-Masanell and Hervas-

providers' promises to protect users' privacy and data security concerns. This is a problem of *accountability*, 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 corresponds to a high level of data security, which is not only implemented by a cloud service provider but also trusted by users. As providers have significantly more technical knowledge and industry experience than most users, the underlying economic problem is *asymmetric information.*

As a consequence of asymmetric information, a leading industry representative summarized the three key problems of the cloud computing industry:[5] First, a standard cloud-specific definition of security is needed. Second, a streamlined process for evaluating cloud service providers is needed. Third, security fears to cloud adoption have to be overcome. The most advanced idea to tackle these issues is "[t]he emergence of a generally accepted cloud security "seal of approval" [that] should allay many of the concerns that stand in the way of this adoption."

Hence, I study the following research questions. How can we incentivize cloud service providers to actually produce high accountability (or data security) levels and to keep their contractual obligations by implementing corresponding (costly) procedures? 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? Could the asymmetric information problem between cloud service providers and users be solved by introducing some type of certification scheme? If so, how to avoid that the certifier hands out certificates to every (paying) provider?

I propose answers to these questions by designing an institution that attenuates the problems of the cloud computing industry and show under which circumstances this institution can exist in equilibrium, thereby improving welfare. The characteristics of the proposed institution are compared with an actually existing one. I construct a game-theoretic model that builds on the insight by practitioners, that a *two-layered* certification mechanism may be able to solve the problem of dishonest certifiers. This scheme is built around a private nonprofit organization that I call *cloud association*, which is governed by representatives of both providers and users

Drane (2014), Campbell et al. (2015), and De Corniere and De Nijs (2016) for entry points to this highly relevant and topical, yet complementary literature.

[5] `https://blog.cloudsecurityalliance.org/2010/05/17/3-problems-cloud-security-certification-can-solve/`.

and which sources the actual auditing process out to a pool of independent for-profit certifiers.[6]

It is shown how and under which conditions the cloud association 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 the cloud association performs random checks of the auditing procedure ("covert testing") and because users who suffered from low accountability of a certified provider would have an incentive to complain with the association. Following a complaint, the cloud association would automatically 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.

Section 2 provides a discussion of the relevant literature and identifies key characteristics of the optimal cloud governance institution. It also reports on an existing cloud certification mechanism. In Section 3, a model of the market interaction between cloud service providers and users (and third parties supporting their transaction) is presented. The model is analyzed in Section 4. Section 5 considers changes to the results if there is (more) competition at the provider or certifier levels, whereas Section 6 discusses implementation issues and practical implications. Section 7 concludes. The Appendix contains an overview of variables and parameters used, extensions, and formal proofs of the model's results.

2 Institutional choice: Reputation, litigation, or certification?

Cloud computing is a highly innovative industry that is driven by technological progress. As research ranging from new aircraft technology (Forbes and Lederman 2013) to new shipbuilding technology in the European middle ages (Masten and Prüfer 2014) has shown, new technologies may require new institutions to govern an industry's professional relationships. But how can users trust providers' announced accountability levels—and, thereby, privacy and data security

[6] Business associations can perform many other tasks that are valuable for their members and have spillover effects on the economy. See Persico (2015), Larrain and Prüfer (2015), and Prüfer (2016) for details.

promises?

Technological solutions: First, the trust problem in cloud computing cannot be solved (exclusively) with technical solutions, for instance, better cryptography or blockchain technologies. The reason is simply that even if these technologies could completely rule out successful attacks of private or state-sponsored hackers to get access to "secure" data, which seems impossible especially in public clouds, users would have to *believe* in absolute data security before increasing economically relevant usage of the cloud. But as all technological solutions, including blockchain or artificial intelligence, have their own problems shattering users' trust, the solution to our problem must contain a mechanism that gets the economic incentives of relevant parties right.

Limitations of public ordering: In search of the optimal institution that may solve this problem, 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).[7]

Second, as exemplified by strong growth (see the introduction), the cloud computing industry is highly innovative and dynamic. Consequently, technological possibilities constantly change, which requires continuous adaptation of products, services, and business models. By contrast, the administrative processes in 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 possibilities. 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 users expose to the cloud.[8]

[7] Current 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.

[8] Prüfer (2013) presents these arguments in more detail.

As a consequence, the optimal cloud governance institution is not characterized by public ordering (state regulation) but by private ordering, that is, by self-governance without any public authority's involvement in rule making, arbitration, or enforcement. Private ordering benefits from more flexible rules and decision-making procedures than public ordering.[9]

Warranties: Amongst private ordering institutions, one standard way to solve asymmetric information 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 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 to face 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.[10] 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.

[9] See Williamson (1999, 2002) for definitions and discussions of the concepts of public and private ordering and Dixit (2009) for a definition and overview of the concept of economic governance.

[10] Both the industry partners and other computer scientists involved in the EU A4Cloud project confirmed this view (`http://www.a4cloud.eu/consortium`) in personal conversations.

Reputation: Expanding the bilateral transaction between seller and consumer, supported by warranties, to other consumers, 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). For the biggest cloud service providers, who are known by virtually everybody in the industry and, hence, about whose trustworthiness information may travel quickly to many users, a classical reputation-based mechanism may be a reasonably effective disciplining device—provided that a sufficient level of competition exists that allows dissatisfied users to switch suppliers. But parts of the attractiveness of new technologies, including cloud computing, is rooted in the innovativeness of small and medium-sized providers. These myriads of innovators cannot be known by everybody in the industry on a global scale, and if they were known, it would be impossible to credibly distribute information about their trustworthiness in a decentralized fashion. Consequently, decentralized information feedback mechanisms may be unreliable and suffer from omitted or biased feedback (Dellarocas and Wood 2008, Cabral and Hortacsu 2010). And if a central player sets up an exchange platform to facilitate information flows, users may not trust that they obtain complete and accurate information about the trading history of a given seller, as long as the incentives and legal obligations of the party managing the exchange platform are unclear to them.[11]

Litigation: 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). 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, Ganuza et al. 2016). 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.

Certification: This paper shows 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 distribution mechanisms, such as those used on many trading platforms

[11] On 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).

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. I propose that *certification agencies* are part of the best institution to support contract enforcement and to prevent opportunistic behavior by cloud service providers: 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.

Certification schemes reduce information load and 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.

Certifiers' credibility: Next to simplifying information complexity, the critical open issue is a 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 from the certified provider?[12] 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 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 reduces certifiers' incentives to get captured in the first place. Studying credit rating agencies, Mathis et al. (2009) conclude that reputation concerns can mitigate shirking if a sufficiently large fraction of a rating agency's

[12] Edelmann (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.

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.

Information transmission problems: 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. 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 heterogenous 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 (as noted above, this may be different only for the biggest and best known providers). 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,[13] and as is argued for by Stahl and Strausz (2011)—they cannot be expected to value users' interests much.

Two-layered certification: 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 take the conjecture of practitioners serious, that a two-layered certification scheme may be able to solve the problem of captured certifiers, and make it a central characteristic of the proposed cloud association institution.

[13] See `https://ic.fsc.org/index.htm`.

A case in point: This coincides with the setup of an actual organization in the cloud computing industry: 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.”[14] Membership for corporate members, usually cloud service providers and corporate cloud customers, is available for U$ 10,000 p.a. and gives members voice within the CSA. A key task of the CSA is to define provider certification schemes and auditor licensing schemes. Under their certification rules, only nationally licensed certification bodies can become certification auditors—provided they become CSA Corporate Members, obtain special training, and pass an exam. The CSA decides when to review the certification bodies, a process that is paid for by the certifiers. It is governed by a board, which hires industry experts as executive managers.[15] The model below reflects this but differs from the CSA in some important aspects, which are discussed in Section 6.

3 A Model of Cloud Governance

Here the model is presented. Subsection 3.2 contains a discussion of key assumptions.

3.1 Model setup

The model features four types of players: users, a cloud service provider, a certification agency, and the cloud association, a nonprofit organization governed by its board. The players play an infinitely repeated game with common discount factor $\delta \in (0, 1)$.

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 in each period $t \in \{1, 2, ..., \infty\}$. Without consumption, a user gets zero value. If user i buys cloud services, at period t she obtains expected consumption utility of:

$$v_i(s_i) = u - (1 - \tilde{\sigma})s_i - p, \tag{1}$$

where $u > 0$ denotes the gross utility from consuming a secure service, and $\tilde{\sigma}$ denotes the probability that no privacy breach (for individual users) or data security incident (for business users) occurs, as expected by users. Henceforth, we will use these terms interchangeably. $\sigma \in$

[14] See `https://cloudsecurityalliance.org/about/`.

[15] See `https://cloudsecurityalliance.org/about/board-of-directors/`.

$[0, 1]$ is a measure for the accountability level set by the provider selling the cloud service, which is unobservable to users. $s_i \sim U(0, \bar{s})$ denotes the disutility of user i in case of a security incident, which captures heterogeneous preferences for accountability across users. 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 denotes the price that a provider charges for one unit of her cloud services.

To increase tractability (by reducing the number of cases when solving the model) and to focus on the interesting case where at least some consumers have significant demand for privacy or data security, we only consider constellations where the following assumption holds.[16]

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

Moreover, to account for the scattered nature of users in the global cloud computing market, we assume that users cannot communicate with each other.

On the supply side, we focus on a monopolistic provider, who maximizes profits and solves, in each period t:[17]

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

where $q(p, \tilde{\sigma})$ 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:

$$c(0) = c'(0) = 0, c'(\sigma) > 0, c''(\sigma) > 0 \quad \forall \sigma > 0. \tag{3}$$

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 can become a member of this association for a fee $F > 0$. In exchange, she can 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. As user representatives, think of industry experts with high personal or

[16] Tsai et al. (2011) show experimentally that users pay a premium to privacy-protective online retailers. Dengler and Prüfer (2018) show theoretically that it is rational for some users to act as if they would value privacy even if don't value it per se.

[17] In Section 5, we discuss qualitative changes if we consider competition among several providers.

institutional reputation for acting in users' interest, for instance, directors of consumer protection agencies or data protection agencies. For simplicity, we assume that the representatives of providers have formal authority over decision making but that users' representatives can veto them.[18] Hence, all decisions issued by the association's board must be acceptable to both users and providers, by construction.

The association's board, however, is limited in size and capacity and, therefore, cannot audit and certify all providers itself. Therefore, the auditing procedure is outsourced to a profit-maximizing *certification agency* (or *auditor*). 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. 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. Seeking 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. 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. Using its 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 ($1 - \sigma$) uncovered by the contract between the provider and a user.

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. If a user complains that $\sigma < \hat{\sigma}$, the association automatically starts an investigation. For 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

[18] If both groups had active decision making power, we would need an assumption about the distribution of bargaining power. This would not affect the results qualitatively but make the solution less tractable.

perfectly detected. 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.[19]

The cloud association also conducts random verification checks of the certifier's audits with probability τ, similarly costing c_V each (τ is the "covert-testing rate"). Think of investigations of the auditing protocol that the certifier has to deliver to the cloud association as part of the standard certification procedure. If the cloud association finds proof of fraud, it revokes the certifier's license and rectifies the provider's certification status.

For tractability, we assume that complaining is costly but not very much so.

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

Cheating of certifiers is additionally prevented by specifying that there are no direct payments 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.

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 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 random certifier among its pool of certifiers with *clean record* (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 to require being audited (for fee F).

3. If an audit is demanded, the **association** commissions a *licensed* certifier to audit the provider, for revenue R. It also sets and announces the "covert-testing rate" τ, with which it randomly tests auditing procedures and certification decisions (for cost c_V each).

[19] This assumption makes sure that the certifier does not always deny certification, independent of true σ.

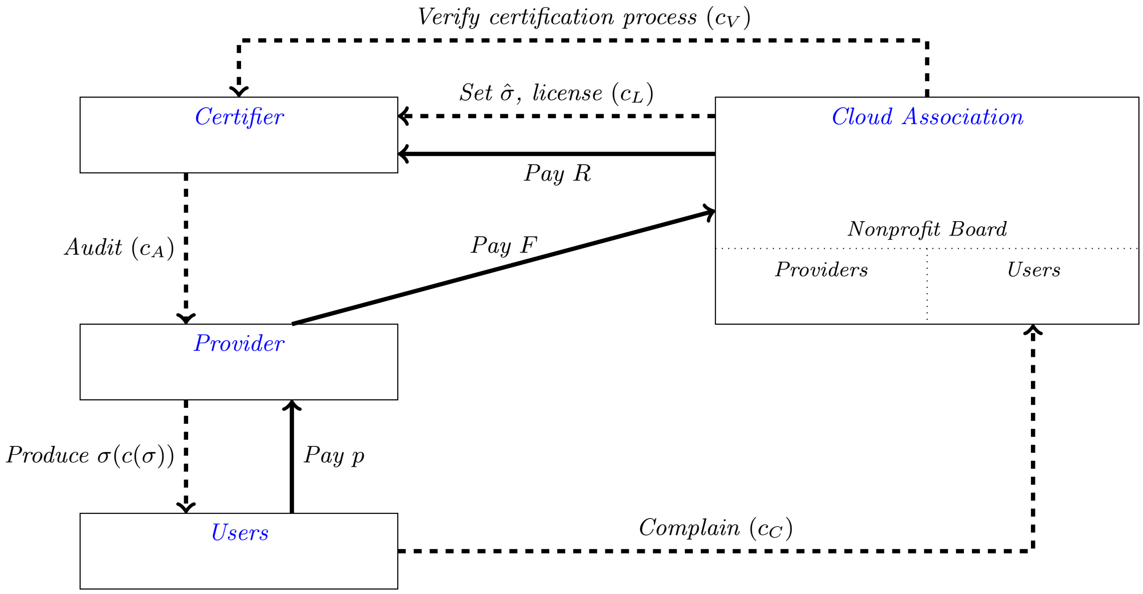

Figure 1: Governing the Cloud with the Cloud Association

4. The commissioned **certifier** chooses the 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.

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 $\sigma < \hat{\sigma}$ (for cost c_C). All complaints are investigated by the cloud association (for cost c_V). If the commissioned certifier was found to have falsely certified a provider, it is deleted from the association's pool of certifiers with clean record.

In period 1, certifiers start with a clean record. As every infinitely repeated game, this game has multiple equilibria. I solve it for a *Markov-perfect equilibrium* that produces $\sigma > 0$ and relies, wherever possible, on pure strategies. The status of a certifier in the cloud association's pool (*with* or *without clean record*) is used as *state variable*, which connects periods t and $t + 1$ and thereby influences a certifier's trade-off at stage 4. All other decisions in period t depend only on information available to the players in period t.

3.2 Discussion of the model

Users' utility function: Assuming positive and heterogeneous s_i is justified because both Chellappa and Sin (2005) and Solove (2007) provide evidence that users have diverse privacy preferences. Goldfarb and Tucker (2012) provide evidence for the fact that users' privacy concerns have been increasing over time. u is constant in equation (1) because independent of the amount of 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 $(1 - \tilde{\sigma})$, scaling a user's expected loss from privacy infringements.

Monopolistic provider: This assumption is not too far from reality as some industry experts predict that Amazon, Google, and Microsoft will capture 76% of all cloud platform revenue in 2018, expanding to 80% by 2020.[20] Nevertheless, assume there is a large set of horizontally-differentiated cloud service providers who could serve the unit mass of atomistic users. Users cannot communicate with each other (because they are scattered globally). Hence, only a bilateral relational contract between one provider and one user would be possible. But the threat of taking away the business of one atomistic user in case of provider shirking (i.e. producing $\sigma < \hat{\sigma}$) is too small to incentivize the provider to invest $c(\sigma) > 0$. Hence, the bilateral relational contract would collapse. In this setup, even if providers compete in prices and certification status, because of horizontal differentiation providers have market power and can, hence, set prices above marginal cost $c(\sigma)$. But providers face a commitment problem, which only the cloud association, serving as a credible central repository of information about provider behavior, can mitigate. As the focus of this paper is to show that the designed cloud association mechanism can produce an equilibrium with $\sigma > 0$, the same situation can be reached by simply assuming that each user faces a monopolistic provider in each period right away.

Solution concept: To avoid provider shirking the certifier can deny to award the certificate at stage 4 of period t, which would hurt the provider at stage 5. But to avoid certifier shirking, the cloud association relies on damaged users complaining at stage 6 (or on high τ). As no punishment by the association is possible in the same period after stage 6, certifier shirking can only be prevented by the shadow of the future: making employment of the certifier in periods $t+1, t+2$, etc. (and constant reimbursement stream R per period) dependent on having a clean record at the end of t. This explains the use of a repeated game and Markov-perfect equilibrium as solution concept.

Incentives of cloud association boards members: 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

[20] `http://www.forbes.com/sites/louiscolumbus/2017/11/07/forresters-10-cloud-computing-predictions-for-2018/`.

providers' (2) interests, respectively.[21] In the model, trustworthiness of board members is merely assumed. In practice, board members must be chosen such that being corrupted is not worth it for them because the economic, legal, and social effects of lost individual reputation in case their shirking becomes publicly known would outweigh the benefits, in expected terms. Think of the incentives of judges or expert witnesses testifying under oath, which are structurally similar to incentives of industry experts representing consumer interests.

Covert testing/random checks of certifier's audits: For monitoring purposes, the cloud association randomly approaches the provider and artificially reduces the provider's accountability level such that $\sigma < \hat{\sigma}$. If, and only if, the certifier invests auditing effort $a = 1$, she can detect these bugs and refuse a certificate to the apparent low-accountability provider. The cloud association would then reveal that these bugs were purposefully prepared, as a sort of "covert testing," delete the bugs, and award the provider with a certificate of its true $\sigma \geq \hat{\sigma}$. By contrast, if the certifier does award the certificate despite the bugs, the cloud association would interpret it as fraud and expel the certifier from its pool of certifiers with a clean record.[22]

Residual risk and user understanding: I assume that a user can distinguish between a privacy infringement because of provider shirking ($\sigma < \hat{\sigma}$) and the residual risk inherent in imperfect security technologies ($1 - \hat{\sigma}$). This is a modeling shortcut. If a user could only see that a security incident occurred but not understand why, she would have to form a belief that the incident occurred because of provider shirking. In Appendix A.2, such a more elaborate Bayesian model is discussed. There I conclude, however, that the additional insights would be limited as compared to the additional model complexity. Therefore, in the model I assume that users suffering from a privacy infringement understand the source of the problem.

[21] Filistrucchi and Prüfer (2018) provide theory and evidence that nonprofit board members make decisions in line with the observable stated values of their employing organizations.

[22] This strategy is reminiscent of Rahman (2012), who studies situations where "monitoring the monitor" is relevant in general. The intuition of his result is provided by examples such as covert testing of airport security officials: the incentives of monitors are kept high by the supervising authority via randomly exposing monitors to "bad" travelers (carrying illegal substances or weapons), which they can only detect when exerting true (unobservable) effort. Importantly, monitors must know that it is possible that such tests take place and that their task is to identify "bad" types. Rahman shows that this strategy leads to high detection rates of truly bad types. See Probst et al. (2012) for more details regarding the technical requirements of ongoing or covert checks.

4 Analysis

The six-stage game in period t is solved by backward induction. I first formally solve the model. Interpretation and economic intuition of the intermediate results are provided after the main proposition. Section 4.8 contains a numerical example and a figure, for illustration. Before the analysis, the incentives of the provider and the resulting expectations of users are clarified.

4.1 Benchmark results

Call the accountability level announced by the provider σ' and the accountability level expected by a user $\tilde{\sigma}$. As a benchmark, we obtain Lemma 1, which is proven in the Appendix.

Lemma 1 (Equilibrium without Certification) *If no certification procedure is in place and users and a provider act in isolation, the unique equilibrium in every period t is s.t. the provider sets $\sigma^* = 0$ and price $p^* = \frac{u}{2}$, which results in equilibrium demand of $q(p^*) = \frac{u}{2\bar{s}}$.*

As second benchmark case, consider a setup where the provider can be certified but where there is no certifier of certifiers. Because users cannot communicate with each other, if a user was shirked by a certified provider, other users could not learn about it and, hence, could not make their strategies dependent on respective information. Consequently, everybody understands that a certificate does not contain valuable information and ignores it. The certifier's best response is to save on auditing effort cost. Lemma 2, mirroring Lemma 1, follows without formal proof.

Lemma 2 (Equilibrium with Certification by an Unmonitored Auditor) *If the auditor is not monitored, the equilibrium in every period t is s.t. the auditor invests effort $a^* = 0$ and hands out a certificate to the provider, or not. The provider sets $\sigma^* = 0$ and price $p^* = \frac{u}{2}$, which results in equilibrium demand of $q(p^*) = \frac{u}{2\bar{s}}$.*

We have seen that a crucial assumption is whether users can (credibly) communicate with each other, or not. Without credible decentralized information, there may be a role for a centralized collector and distributor of information about the provider's and a certifier's conduct. Assume, for now, that such a credible central player is in place who hands out certificates veritably, that is, if, and only if, the provider's $\sigma \geq \hat{\sigma}$. It is shown below under which circumstances this constitutes an equilibrium.

In this case, excess accountability above $\hat{\sigma}$ is still unobservable and hence would not be rewarded by users by paying a higher price. 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, drawing on Lemma 1:

$$\tilde{\sigma} = \begin{cases} \hat{\sigma} & \text{if } x = 1, \\ 0 & \text{if } x = 0. \end{cases} \tag{4}$$

4.2 Stage 6

Having established benchmark results, I now solve the full game as described in Section 3. 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:

$$\begin{aligned} u - (1 - \hat{\sigma})s_i - c_C &\geq u - (1 - \sigma)s_i & (5) \\ \Leftrightarrow \quad c_C &\leq (\hat{\sigma} - \sigma)s_i & (6) \end{aligned}$$

Using Lemma 1, which suggests that $\sigma = 0$ if $\sigma < \hat{\sigma}$, we find:

Lemma 3 (Equilibrium user complaints) *(i) If, and only if, the cost of complaining is sufficiently small, cheated users will complain to the association (if $c_C \leq \hat{\sigma} s_i \equiv \overline{c_C}$). (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.[23] Consequently, users know that, even if they get a low-accountability service from a certified provider, they will receive a substitute service compensating their damage immediately (and, in the case of a second low-accountability service, receive another substitute, etc.). Thus, they can trust the certificate and ignore potential deviations of providers when making consumption decisions.

4.3 Stage 5

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 (4), it follows that all users with

[23] See Appendix A.2 for the adapted results and a simple solution if complaining is costly.

accountability preferences

$$s_i \leq \frac{u - p(x=1)}{1-\hat{\sigma}} \equiv \overline{s_1} \tag{7}$$

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

$$s_i \leq u - p(x=0) \equiv \overline{s_0} \tag{8}$$

buy cloud services if $x = 0$, where $p(x = 1)$ and $p(x = 0)$ are determined at stage 2.

4.4 Stage 4

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 saves c_A this period but can only make a guess about the provider's true σ. If the certifier would not spend c_A but still award a certificate, the provider might set $\sigma = 0$ with positive probability. In this case, however, because of Lemma 3, the fraud would be detected, the certifier would lose her clean record and, hence, never obtain an auditing job in the future again.[24]

Moreover, the cloud association has announced to check every award of a certificate with probability τ. It can do this—and detect fraud with certainty—because the certifier is required to use the cloud association's certification protocol after being licensed at stage 1. The cloud association's announcement is credible because its board is comprised of representatives of both market sides, users and providers, who understand that random checks are crucial to sustain the cloud association certification system. Therefore, if in period t the certifier chooses $a = 1$, she expects a net present value of $\frac{1}{1-\delta}(R - c_A)$ from future business with the cloud association. If she chooses $a = 0$ (but returns to $a = 1$ thereafter, according to the one-stage deviation principle), she expects a net present value of, $R + (1-\tau)\frac{\delta}{1-\delta}(R - c_A) + \tau * 0$. Hence, the certifier prefers $a = 1$ over $a = 0$ if, and only if:

$$\begin{aligned} \frac{1}{1-\delta}(R - c_A) &\geq R + (1-\tau)\frac{\delta}{1-\delta}(R - c_A) && (9) \\ \Leftrightarrow \quad \tau &\geq \frac{(1-\delta)c_A}{\delta(R - c_A)} \equiv \underline{\tau}(R) && (10) \end{aligned}$$

[24] If 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, also implies a loss of all future auditing mandates for the certifier.

The left-hand side (LHS) of (9) 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, thereby saving c_A, and earning $(R - c_A)$ in the future if the cloud association detected no shirking. Re-arranging yields (10), which specifies the required minimum rate of random checks, $\underline{\tau}(R)$, that deters cheating by the certifier.

4.5 Stage 3

At stage 3, the cloud association commissions the licensed certifier and determines the revenue R it offers for conducting the audit. It also determines τ, the rate of random verification checks or "covert testing" of the certifier's audits, for unit cost c_V. Solving (9) for R specifies the revenue incentivizing the certifier to choose $a = 1$:

$$\underline{R} = \left(\frac{(1-\delta)}{\delta\tau} + 1 \right) c_A \tag{11}$$

Note that $\frac{\partial R}{\partial \tau} < 0$. This implies that the cloud association has two instruments, R and τ to discipline the certifier. It can either pay a high revenue for each audit until eternity—which the certifier does not want to risk by shirking on auditing effort—or it can increase the rate of random checks.

(11) specifies the direct cost of one certification, whereas τc_V captures the associated expected verification cost. In order to minimize its total cost, the cloud association solves:

$$min_\tau \quad \left(\frac{(1-\delta)}{\delta\tau} + 1 \right) c_A + \tau c_V \tag{12}$$

Solving (12) for τ and substituting the result into (11) yields the following Lemma.

Lemma 4 (Equilibrium certifier revenues, covert testing, and auditing decisions) *(i) The cloud association offers a certifier revenue $R^* = c_A + \sqrt{\frac{(1-\delta)c_A c_V}{\delta}}$ per audit and randomly verifies every audit with probability $\tau^* = \sqrt{\frac{(1-\delta)c_A}{\delta c_V}}$. Its total expected cost per certification are therefore $c_A + 2\sqrt{\frac{(1-\delta)c_A c_V}{\delta}}$. (ii) The certifier accepts the auditing job, invests effort $a = 1$, and makes truthful certification decisions.*

This Lemma contains several insights. First, by trading off the direct cost per certification procedure and the indirect cost by ex post verification or covert testing, we find a unique tuple of revenues and random checks, (R^*, τ^*), that both incentivizes the certifier to invest auditing effort and that minimizes the total costs that the cloud association has to bear per certification procedure. These costs increase in c_A and c_V (the procedural costs of auditing with effort and verifying the auditing procedure) and they decrease in δ, the discount factor of the certifier.

Notably, $R^* > c_A$: the certifier earns an information premium at equilibrium and makes positive expected profits every period she is commissioned to audit the provider. Because the certifier fears losing this rent in the future if shirking now, she invests effort in auditing.

4.6 Stage 2

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 (7) and (8), the provider faces demand $q(x=1) = \overline{s_1}/\bar{s}$ or $q(x=0) = \overline{s_0}/\bar{s}$, depending on her certification status. Given (4), seeking certification implies setting $\sigma = \hat{\sigma}$ and not seeking certification implies $\sigma = 0$. In the appendix, we prove the following Lemma.

Lemma 5 (Provider incentives and pricing) *The provider joins the association for the fee F and seeks certification if, and only if:*

$$\hat{\sigma} \geq \begin{cases} 1 - \frac{u}{2\bar{s}} & \textit{if } u \in (0, \frac{1}{2}(\bar{s} - \sqrt{5\bar{s}^2 - 8\bar{s}(c(\sigma)+F)})) \\ \frac{(\bar{s}-u)^2}{2\bar{s}^2} + \frac{c(\sigma)+F}{\bar{s}} & \textit{if } u \in [\frac{1}{2}(\bar{s} - \sqrt{5\bar{s}^2 - 8\bar{s}(c(\sigma)+F)}, 2\bar{s})) \end{cases} \quad \textit{and } c(\sigma) + F < \frac{5}{8}\bar{s} \tag{13}$$

or if

$$\hat{\sigma} < 1 - \frac{u}{2\bar{s}} \quad \textit{and } c(\sigma) + F < \frac{u}{2} + \frac{u^2}{4\bar{s}}. \tag{14}$$

If the provider seeks certification, she sets $\sigma = \hat{\sigma}$ and prices cloud services as follows:

$$p^* = \begin{cases} u - (1-\hat{\sigma})\bar{s} & \textit{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}, \\ \frac{u}{2} & \textit{if } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \end{cases} \tag{15}$$

If she does not seek certification, she sets $\sigma = 0$ and $p^* = \frac{u}{2}$.

Lemma 5 shows that the provider can be incentivized both for high or for low required accountability levels $\hat{\sigma}$ to join the association and to set $\sigma = \hat{\sigma}$. Whether this occurs at equilibrium depends on parameter realizations. In particular, the provider's costs to produce accountability ($c(\sigma)$) and to join the cloud association (F) must not be too high.

4.7 Stage 1

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 and covert testing of one provider, namely R^*, c_L, and c_V. Using Lemma 4.(i), adding c_L, and realizing that there is no gain if the (nonprofit) cloud association piles up cash, the membership-fee is $c_L + c_A + 2\sqrt{\frac{(1-\delta)c_A c_V}{\delta}}$.

Regarding $\hat{\sigma}$, the preferences of the board representatives of providers and users are different, however. Drawing on A.7, the representatives of providers solve the following program:

$$max_{\hat{\sigma}} \quad \pi = \begin{cases} u - (1-\hat{\sigma})\bar{s} - c(\sigma) - F^* & \text{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}, \\ \frac{u^2}{4(1-\hat{\sigma})\bar{s}} - c(\hat{\sigma}) - F^* & \text{if } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}. \end{cases} \tag{16}$$

$$\textit{s.t. (13) and (14) hold and } \pi(\hat{\sigma}) \geq 0.$$

The side-constraints (13) and (14) make sure that seeking certification is incentive compatible for the provider; $\pi(\hat{\sigma}) \geq 0$ makes sure that the provider's participation constraint holds. The upper row of (16) refers to the covered market case, which is triggered by high $\hat{\sigma}$, the lower row captures the case where some consumers do not buy. I prove Lemma 6 in the appendix.

Lemma 6 (Equilibrium membership-fee and accountability preferences) *(i) The equilibrium membership-fee for each provider who seeks being audited is* $F^* = c_L + c_A + 2\sqrt{\frac{(1-\delta)c_A c_V}{\delta}}$. *(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:*

$$\hat{\sigma}_P^* = \begin{cases} \hat{\sigma}_1 & \textit{if } \frac{u^2}{4(1-\hat{\sigma}_1)\bar{s}} - c(\hat{\sigma}_1) > u - (1-\hat{\sigma}_2)\bar{s} - c(\hat{\sigma}_2), \\ \hat{\sigma}_2 & \textit{otherwise} \end{cases} \tag{17}$$

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 (7) we derive users' equilibrium consumption decisions (at stage 5) and the associated payoffs.

Lemma 7 (Equilibrium consumption and payoffs) *(i) In the certification equilibrium, all users with accountability preferences* $s_i \leq \overline{s_1}^*$ *buy the provider's service for the price* p^*, *where* $\overline{s_1}^* = \frac{u}{2(1-\hat{\sigma}^*)}$ *if* $c'(\hat{\sigma}^*) < \bar{s}$, *and* $\overline{s_1}^* = \bar{s}$ *if* $c'(\hat{\sigma}^*) \geq \bar{s}$. *(ii) Consumer surplus amounts to:*

$$CS = \begin{cases} \frac{(2\bar{s}-1)u^2}{8\bar{s}^2(1-\hat{\sigma}^*)} & \textit{if } c'(\hat{\sigma}^*) < \bar{s}, \\ \frac{1}{2}(2\bar{s}-1)(1-\hat{\sigma}^*) & \textit{if } c'(\hat{\sigma}^*) \geq \bar{s}. \end{cases} \tag{18}$$

(iii) The provider obtains surplus of:

$$\pi^* = \begin{cases} \frac{u^2}{4(1-\hat{\sigma}^*)\bar{s}} - c(\hat{\sigma}^*) - (\frac{c_A}{\delta} + c_L + c_V) & \textit{if } c'(\hat{\sigma}^*) < \bar{s}, \\ u - (1-\hat{\sigma}^*)\bar{s} - c(\hat{\sigma}^*) - (\frac{c_A}{\delta} + c_L + c_V) & \textit{if } c'(\hat{\sigma}^*) \geq \bar{s}. \end{cases} \tag{19}$$

(iv) Total welfare is given by:

$$W^* = \begin{cases} \frac{(4\bar{s}-1)u^2}{8\bar{s}^2(1-\hat{\sigma}^*)} - c(\hat{\sigma}^*) - (\frac{c_A}{\delta} + c_L + c_V) & \textit{if } c'(\hat{\sigma}^*) < \bar{s}, \\ u - \frac{(1-\hat{\sigma}^*)}{2} - c(\hat{\sigma}^*) - (\frac{c_A}{\delta} + c_L + c_V) & \textit{if } c'(\hat{\sigma}^*) \geq \bar{s}. \end{cases} \tag{20}$$

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

We summarize equilibrium decisions in the following Proposition, the model's key result.

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

1. *The cloud association sets a membership-fee $F^* = c_L + c_A + 2\sqrt{\frac{(1-\delta)c_A c_V}{\delta}}$ and a threshold accountability level:*

$$\hat{\sigma}^* = \begin{cases} \hat{\sigma}_1 & \textit{if } \frac{u^2}{4(1-\hat{\sigma}_1)\bar{s}} - c(\hat{\sigma}_1) > u - (1-\hat{\sigma}_2)\bar{s} - c(\hat{\sigma}_2), \\ \hat{\sigma}_2 & \textit{otherwise} \end{cases} \tag{21}$$

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 price:*

$$p^* = \begin{cases} \frac{u}{2} & \textit{if } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \\ u - (1-\hat{\sigma})\bar{s} & \textit{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}. \end{cases} \tag{22}$$

3. *The association commissions a certifier and offers revenue $R^* = c_A + \sqrt{\frac{(1-\delta)c_A c_V}{\delta}}$ and randomly verfies every audit with probability $\tau^* = \sqrt{\frac{(1-\delta)c_A}{\delta c_V}}$.*

4. *The certifier sets $a^* = 1$ and awards a certificate truthfully. The provider does not complain.*

5. *Users with accountability preferences $s_i \leq \overline{s_1}^*$ buy from the certified provider; users with preferences $s_i > \overline{s_1}^*$ do not buy. $\overline{s_1}^* = \frac{u}{2(1-\hat{\sigma}^*)}$ if $c'(\hat{\sigma}^*) < \bar{s}$, and $\overline{s_1}^* = \bar{s}$ 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 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 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" makes sure that the association has funds to pay for the ex ante licensing of the auditor, that it can pay the auditor an efficiency wage, which restores her incentives to work and be honest when auditing, and that can also fund the cost of random checks of the auditor's work.[25]

This efficiency wage, R^*, grows in the cost the auditor incurs when exerting effort (c_A), in the cost of ex post verification (c_V), and if the auditor is less patient (if δ decreases). Moreover, the cloud association chooses the tuple (R^*, τ^*) s.t. total expected costs of certification are minimal: from the auditor's perspective, a higher revenue R has the same incentive effect to award a certificate truthfully as a higher covert-testing rate τ. Thereby, the auditor makes positive equilibrium profits—she earns an information rent—but behaves as aspired by the association because she fears losing those profits in the future when her fraud 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 first insight 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 $\hat{\sigma}^*$. Its 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 could not 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 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. 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 (22) states 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 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

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

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 $\bar{s}$ (see Section 4.8 for illustration).

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 about receiving low accountability is rather low. In this baseline model, Assumption 2 does the trick for us, assuming $c_C \to 0$. In Section 5, 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 an 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 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 in the subgame after provider's defection. If this holds, however, actual punishment costs are saved. Here the saved costs on the equilibrium path are the cost of complaining (c_C) and, partly, the cost of ex post verification of the auditing procedure ($(1-\tau)c_V$). The costs actually incurred on the equilibrium path are the cost of licensing certifiers (c_L), of auditing providers (c_A), and of random verification checks of the certifier's work ((τc_V).

Further insights are generated by studying equilibrium payoffs (see Lemma 7). 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 privacy 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. Because users are represented on the cloud association board and equipped with the right to veto the strategy proposals of providers, 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. This characteristic survives when there are two competing providers (see section

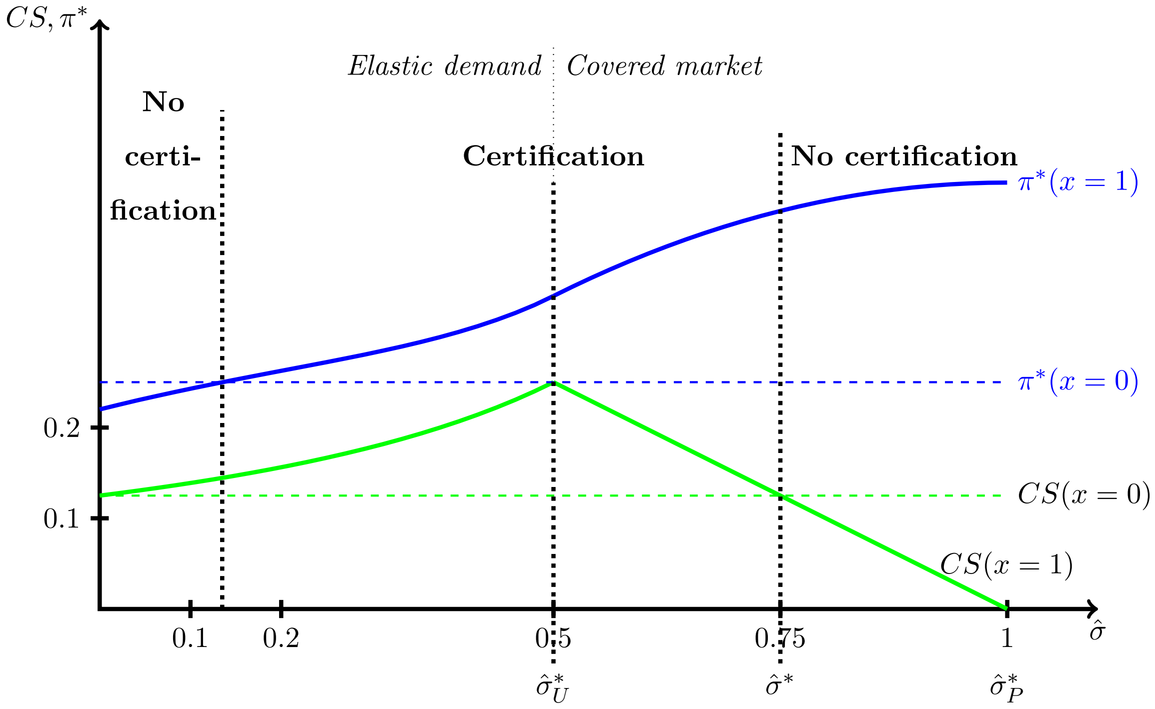

Figure 2: Consumer surplus and profits with varying $\hat{\sigma}$ (numerical example).

5.1). Consequently, this model predicts that it will be providers who stop using the certification scheme first if the usage costs grow too large.

4.8 Numerical example

For illustration, consider the following numerical example, which is displayed in Figure 2. Assume $u = \bar{s} = 1$, which satisfies Assumption 1. Assume $c_L + c_A + 2\sqrt{\frac{(1-\delta)c_A c_V}{\delta}} = 0.03 = F^*$, and $c(\sigma) = \frac{\sigma^2}{2}$, which satisfies conditions in equation (3).

In this example, demand is elastic for all $\hat{\sigma} < 1 - \frac{u}{2\bar{s}} = \frac{1}{2}$. Drawing on Lemma 1, certification is attractive for the provider if (19)$> \frac{u^2}{4\bar{s}} = 0.25$. For the elastic case, (A.9) holds if $\hat{\sigma} > 0.135$. For the inelastic case, (A.8) holds for all supported values. Profits are monotonically increasing in $\hat{\sigma}$. The global profit maximum is at $\hat{\sigma}_P^* = 1$. Substituting numbers in (A.17) and (A.18) shows that consumers prefer buying services from a certified provider over a non-certified provider for all $\hat{\sigma} > 0$ in the elastic case and for $\hat{\sigma} \leq 0.75$ in the inelastic case. Consumer surplus is maximized at $\hat{\sigma}_U^* = 1 - \frac{u}{2\bar{s}} = 0.5$. As provider representatives want to approach $\hat{\sigma}_P^* = 1$ as much as possible, the inelastic case is relevant. Consumer representatives veto $\hat{\sigma}$-levels larger than 0.75 because there consumer surplus is higher without certification. Hence, the equilibrium accountability level set by the cloud association is $\hat{\sigma}^* = 0.75$.

In Figure 2, the solid curves denote consumer surplus and provider's profits with certification

($x = 1$), whereas dashed horizontal lines denote consumer surplus and provider's profits without certification ($x = 0$), all as functions of the threshold accountability level, $\hat{\sigma}$. Certification is only an equilibrium if both groups (weakly) benefit from it. For very low levels of $\hat{\sigma}$, certification is unprofitable for providers as the amount of additional users willing to pay for certified cloud services is limited and providers are better off by offering zero accountability for a low (but positive) price. Very high levels of $\hat{\sigma}$, by contrast, are nice to have in principle—but due the convex cost of accountability providers then have to charge prices that cannot be paid by sufficient amounts of consumers anymore.

5 Competition

5.1 Competing cloud service providers

In the baseline model we have taken a shortcut by assuming a monopolistic provider. This has kept the model tractable. A more elaborate competition model would have to differentiate among several cases (not only elastic demand versus covered market) and hence could not include the entire spectrum of players we deem necessary for the proposed institution (cloud association, certifiers, provider, users). But if we were to model competition in a less stylized way, what type of model would be appropriate? Classical Cournot, Bertrand, or Hotelling (or Salop-circle) models endogenize competitors' quantity or pricing or location decisions. But quantity (or capacity) decisions are decidedly not the key issue in cloud computing because one feature of cloud business models is that additional storage or computing facilities are available at very short notice.[26] This even counts for potential capacity constraints of cloud service providers themselves, which can be alleviated by sourcing from other providers. The resulting issue of certification of an entire cloud service supply chain is discussed in Appendix A.3.

By contrast, the key endogenous variable in my model is accountability, which relates to ex ante non-verifiable quality. The closest classical competition model to capture this is the Shaked and Sutton (1982) model of vertically differentiated duopoly producers. Shaked and Sutton's key result is that two producers who simultaneously first set quality levels and then prices will vertically differentiate their products in equilibrium. Translated to cloud computing, if two cloud service providers simultaneously play the game constructed in Section 3, 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

[26] See the definition of cloud computing in the first line of p.1 of this paper or `https://datacenterfrontier.com/capacity-constraints-how-prepare-for-future/` for an industry blog confirming this view.

join the association, produce no accountability ($\sigma = 0$) and charge a low but strictly positive price (see Lemma 1). Consequently, just as in Proposition 1, users with high accountability preferences $s_i > \overline{s_1}^*$ do not buy (potentially an empty set). In the duopoly, however, cloud services are vertically differentiated and, hence, the market is split: users with intermediate preferences, $s_i \in (\overline{s_0}^*, \overline{s_1}^*]$, buy from the certified provider, whereas users with low valuation for accountability, $s_i \leq \overline{s_0}^*$, buy from the uncertified provider; where $\overline{s_0}^* < \overline{s_1}^*$.

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 significant 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. 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 accountability that a certified provider can credibly supply, as compared to the uncertified outside option, $\hat{\sigma}$, 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 larger than 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, net of market entry costs, are zero. In the model set-up of this paper, given that the cost for uncertified providers is normalized to zero, many firms could offer uncertified services and de facto find themselves in a market with Bertrand price competition with homogeneous 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 $(\overline{s_1}^* - \overline{s_0}^*)$ is split among those providers, which reduces revenues of each individual provider. Such entry could be incentive-compatible for a limited number of certified providers at equilibrium, given the right parameter values. Once demand is split too much due to entry, 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 to introduce 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 the motivation 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.

5.2 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 incentive compatibility constraint of the certifier, the equivalent of (9), assuming that the certifier has exactly one auditing job in period t, would turn into:

$$(R - c_A) + \frac{\delta}{1-\delta}\frac{M}{N}(R - c_A) \geq R + (1-\tau)\frac{\delta}{1-\delta}\frac{M}{N}(R - c_A) \tag{23}$$

This would affect both the equilibrium levels of R^* and τ^* but would not change the analysis or the quality of the results, apart from the limits discussed in section 5.1.

6 Implementation Challenges

The cloud association governance mechanism proposed here is novel. But related institutions do exist in practice, as the description of the Cloud Security Alliance (CSA) in Section 2 has shown. The key differences between the existing CSA scheme and the institution proposed in this paper are the following: First, my proposed scheme does not rely on governmental pre-licensing of certifiers by their national governments (because of the problems of public ordering discussed in 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 allows 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. Although the adaptation of all three differences is important—a point that was acknowledged by a CSA Director in personal communication—the second one is crucial: If they want to increase users' trust in their data-security announcements, providers must realize that it is in their own interest to establish knowledgeable and powerful experts defending users' interests on the association's board. Only then, and if this characteristic is communicated widely, users will rationally trust the certification decisions of providers more, which will lead to more business and higher profits of providers in the end (remember footnotes 2 and 3).

Nevertheless, caution is warranted when implementing the cloud association institution because the model presented here, as any economic model, relies on assumptions that are simplified versions of reality. A key issue, which motivated the two-layered certification structure modeled, regards the threat of bribing certifiers, who may be tempted to falsely hand out certificates to providers in exchange for tangible benefits. Lemma 4 has shown that it is crucial to offer certifiers efficiency wages. This means, to keep certifiers honest they must believe that they can earn a good profit with working for the Cloud Association in the future to prevent them from accepting a bribe. Expecting to earn a "good profit," however, depends on various preconditions. First, competition among certifiers to get an auditing job must not be too intense. For that purpose, the Cloud Association should not enlarge its pool of certifiers that can be called for auditing jobs too much. It should also refrain from the (tempting) option to renegotiate (read: to lower) fees for certifiers as this would destroy the relational contract between a certifier and the association. An obvious proximate risk if the Cloud Association scheme gets very successful but some would-be certifiers are not licensed is that those rebuffed certifiers open a competing association and offer to work for lower fees than in the first one, which would even allow the second association to undercut the certification prices of the first one (see F^* in Lemma 6). To prevent such undercutting, the Cloud Association might want to advertise widely that only an association that keeps certifiers incentivized to act honestly by paying them reasonable fees can credibly uphold high accountablity levels.

Complementing this strategy, it can help to recall Lemma 4, which shows that the fee for certifiers per auditing and the "covert testing" rate of random checks are substitutes: if the market structure is more competitive such that the fee offered to certifiers can or must actually be decreased, the certifier's incentives to act honestly can be upheld be increasing the rate of random checks.

Similarly, the Cloud Association must make sure that once-defaulted certification agencies cannot change their name and pop up as candidate certifiers with a clean record again (*whitewashing*). If this was possible, it would undermine the threat of punishment by foregone business with the association and thereby reduce the incentives of certifiers to not accept bribes if offered by providers. Dealing with such problems is possible, however, by using background checks of the involved persons and assets as part of the licensing procedure that candidate certifiers have to undergo before they are accepted in the pool of certifiers with clean record.[27]

Next to certifiers, the other players in the game must be incentivized to behave as outlined

[27] The *International Cotton Association* offers a great example how default lists are used to identify defectors and to keep them away from profitable transactions. See `http://www.ica-ltd.org/76-arbitrations-106-defaults-30-appeals/` for details.

in Proposition 1. This may be harder for cloud service providers in practice than in theory. For instance, the model assumes that all players know providers' cost function to produce accountability ($c(\sigma)$) precisely. Therefore, equilibrium prices, accountability levels, and demand can all depend on the threshold accountability level ($\hat{\sigma}$), which depends on the cost function in equilibrium. But if costs to produce certain accountability levels are not public information, the other players will have to work with expectations and heuristics, which can lead to errors and destabilize the described equilibrium. Notably, this is also important regarding the details of the complaints mechanism by which consumers who received a low-accountability service from a certified provider can inform the Cloud Association about their problem. Such details are discussed in Appendix A.2.

An especially critical case for implementation are the incentives of users' representatives on the Cloud Association board. If they can be corrupted, the mechanism will break down. In the model description, I suggested to only select individuals with high personal reputation for integrity and for advancing the cause of users, such as the presidents of consumer protection agencies or nongovernmental organizations who have visibility and a verifiable track record.[28] But in practice it may help the system to work if these first-rate individuals representing users' interests are checked (potentially randomly) by outside experts and if they have to be transparent regarding their incomes, in order to compound side-payments from providers who want them, for instance, not to veto certain policy changes on the Cloud Association board.

Complementing these suggestions, despite my advocacy of private ordering to govern the cloud computing industry in Section 2 and the model in Section 3 to be set up without any role for governmental authorities, the system can become more robust if public authorities complement (rather than substitute) private parties. As documented in Prüfer (2016, see esp. footnote 32), industry associations have been supported by the enforcement powers of state authorities since the middle ages, even if their communication and arbitration mechanisms have worked entirely with private actors. As of 2018, the EU General Data Protection Regulation (`https://eugdpr.org/`), a piece of legislation that is enforced by state authorities, is not legally binding outside of the EU—but as many firms from outside of the EU sell services within the EU, it can leverage impact beyond its original jurisdiction. Applying this thought to the case of the proposed Cloud Association, state actors such as tax authorities could play a role in

[28] The empirical findings of Filistrucchi and Prüfer (2018) are encouraging that it is possible to find such individuals. They show that the managers of religious nonprofit hospitals in Germany act in such a consistent way in line with the theological doctrines of their parent churches (that is, with a predefined objective function) that it is possible to differentiate Catholic and Protestant hospitals just by comparing the strategies they use in the market.

confirming that a candidate for a consumer-representative board member position has never taken any bribes or undisclosed income. And law enforcement agencies could help to detect candidate certification agencies with a criminal record, possibly even outside of the jurisdiction in which the Cloud Association is situated.

Moreover, when implementing the proposed scheme 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.

7 Conclusion

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 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. Even if a cloud service provider implemented a high accountability level, she has no means to credibly convince users that their data are secured (as long as no enforcement mechanism relying on third parties, as proposed here, is used and as long as the provider and her customer are not large enough to engage in a bilateral relational contract). As soon as certification agencies, who audit providers' accountability levels and award "trust" certificates, are employed, information problems can be mitigated but an additional problem arises: certifiers' moral hazard to actually spend unverifiable effort on understanding the true accountability levels of providers and not getting captured.

Inspired by existing certification schemes, I develop and apply the idea that representatives of both providers and users (technically-skilled industry experts with high personal reputation for integrity) should jointly oversee the certification process in an independent, private body, the *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. Notably, this mechanism is the simplest possible implementation of a two-tier certification structure and in line with industry practice.

The results of the model 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). 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 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: 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 certifiers' past actions.

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 production and distribution, such that capacity constraints are usually not binding and such that sellers sell their products worldwide to heterogenous 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.

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 in several disciplines, including computer science and law, to study its optimal implementation. Insights from legal scholarship may be helpful in identifying the necessary legal circumstances—even without active involvement of governments—when establishing the cloud association framework, and further research in cryptography and data security infrastructures may impact parameters suggested in the model at hand. 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.

A Appendix

A.1 Overview of variables and parameters used (in order of appearance)

Variable/Parameter	**Definition**
δ	discount factor
t	time index for a period
v_i	indirect consumption utility of user i (net of costs)
u	gross consumption utility
$s_i \sim U(0, \bar{s})$	disutility of user i if privacy breach occurs (“taste for privacy”)
σ	prob that no privacy breach occurs (“accountability”)
p	price charged by provider to user for one unit of cloud services
$\tilde{\sigma}$	accountability level expected by users
$q(p, \tilde{\sigma})$	demand for provider’s services
$c(\sigma)$	cost of producing accountability σ
F	membership fee for provider in cloud association (=cost of being audited)
$\hat{\sigma}$	threshold accountability to become a certified cloud service provider
a	effort level of certifier when auditing a provider
c_A	auditing cost (of certifier)
c_L	cost of licensing (borne by cloud association)
τ	prob of random verification check of an audit (“covert/testing rate”)
c_C	cost of complaining with cloud association about low accountability (for user)
R	remuneration of certifier for performing an audit (from cloud association)

A.2 Rethinking the complaints mechanism

Assumption 2 states that the costs of complaining for users are very low. If users experiencing a privacy breach do not understand when it is due to residual risk $(1 - \hat{\sigma})$, they may complain frequently, each time leading to an investigation of the case by the cloud association and to costs c_V. This could render the entire mechanism prohibitively costly. One response of the cloud association could be to wait until a certain mass of complaints, g, about one provider have piled up before opening an investigation. Denote the share of complaining users of a provider that is necessary to trigger an investigation by $\gamma \equiv \frac{g}{\bar{s_1}^*}$. As consumers are rational in this model, they would form a belief about γ: call ω the probability expected by a user that her complaint would trigger an investigation, where $\frac{\partial \omega}{\partial \gamma} < 0$. Then the user also expects to be reimbursed with probability ω (not 1, as in the baseline model). Manipulating (5) and (6) accordingly shows that a user who received a cloud service with $\sigma < \hat{\sigma}$ will complain if, and only if:

$$u - \omega(1-\hat{\sigma})s_i - (1-\omega)(1-\sigma)s_i - c_C \geq u - (1-\sigma)s_i \tag{A.1}$$

$$\Leftrightarrow \quad c_C \leq \omega(\hat{\sigma} - \sigma)s_i = \omega\hat{\sigma}s_i \tag{A.2}$$

Hence, with lower ω fewer users experiencing low accountability would actually complain. Because $\frac{\partial\omega}{\partial\gamma} < 0$, a higher threshold g decreases users' beliefs that complaining is worth the cost. At a Bayesian equilibrium, we would have that users' belief about $\omega^*(\gamma)$ and the actual threshold $\gamma^*(\omega)$ reinforce each other. These changes would not affect the rest of the model's results qualitatively, however.

As a second alternative specification to the baseline model, consider relaxing Assumption 2. Then a user who complains about low accountability of a provider, bears a substantial cost $c_C > 0$. In this case, Lemma 3.(ii) states that only security-sensitive users (with $s_i \geq \frac{c_C}{\hat{\sigma}}$) would complain in equilibrium. This implies that if a provider produces $\sigma < \hat{\sigma}$, mass $\frac{c_C}{\hat{\sigma}}$ in all $\bar{s}_1^*$ buyers would not complain. If c_C is large, this might invite some providers to 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 (9), would become harder to fulfil because in expectation the certifier would get payoff R in every future period with probability $\frac{c_C}{\hat{\sigma}}$ despite having cheated. This effectively increases the necessary efficiency wage R^* above the level specified in Lemma 4.

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 $\sigma < \hat{\sigma}^*$, independent of their s_i. 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).

A.3 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 (*data controller*), who concludes a service level agreement with the user (*data owner*, in case of individual users potentially also *data subject*), would be responsible for the accountability level of all upstream providers (*data processors*).[29] 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 $k \in \{1, ..., K\}$, at equilibrium we must have $argmin\{\sigma_k \in \{\sigma_1, ..., \sigma_K\}\} = \hat{\sigma}^*$; see equation (21). The weakest link of the chain must be sufficiently strong to sustain credibility in the entire chain.

A.4 Proof of Lemma 1

σ is unobservable to users. Hence, because $c(\sigma)$ is increasing in σ and because the provider has no tool to credibly commit to $\sigma > 0$, she maximizes profits by setting $\sigma = 0$. Users realize this incentive. Consequently, independent of σ', without credible certification we have $\tilde{\sigma} = 0$ at equilibrium, confirming the provider's incentive at equilibrium. In this case, (1) implies a demand function of:

$$q(p) = \begin{cases} 1 & \text{if } p \leq u - \bar{s}, \\ \frac{u-p}{\bar{s}} & \text{if } p \in (u - \bar{s}, u) \\ 0 & \text{if } p \geq u. \end{cases} \tag{A.3}$$

The provider's profit is $\pi(\sigma = 0) = pq(p)$. Using Assumption 1, profit is maximized by setting $p^* = \frac{u}{2}$, which leads to demand of $q(p^*) = \frac{u}{2\bar{s}}$ and provider profits of $\pi^* = \frac{u^2}{4\bar{s}}$. *Q.E.D.*

A.5 Proof of Lemma 5

For $\tilde{\sigma} = 0$, the case without certification analyzed in Lemma 1 applies: at equilibrium $p^* = \frac{u}{2}$, $q(p^*) = \frac{u}{2\bar{s}}$, and $\pi(p^*) = \frac{u^2}{4\bar{s}}$.

[29] Pearson (2011) explains the terminology used in more detail.

For $\tilde{\sigma} = \hat{\sigma}$, (1) implies a demand function of:

$$q(p) = \begin{cases} 1 & \text{if } p \leq u - (1-\hat{\sigma})\bar{s}, \\ \frac{u-p}{\bar{s}} & \text{if } p \in (u - (1-\hat{\sigma})\bar{s}, u) \\ 0 & \text{if } p \geq u. \end{cases} \tag{A.4}$$

The first row of (A.4) captures the case of inelastic demand (or covered market), the second row captures the case of elastic demand, while pricing in the third row is prohibitive. Consequently, the provider's profit function is:

$$max_p \quad \pi = \begin{cases} p - c(\hat{\sigma}) - F & \text{if } p \leq u - (1-\hat{\sigma})\bar{s}, \\ p\frac{u-p}{(1-\hat{\sigma})\bar{s}} - c(\hat{\sigma}) - F & \text{if } p \in (u - (1-\hat{\sigma})\bar{s}, u), \\ -c(\hat{\sigma}) - F & \text{if } p \geq u. \end{cases} \tag{A.5}$$

The third case is clearly dominated. In the first row of (A.5), given that a further price decrease cannot increase demand of a covered market, the profit-maximizing price is $p = u - (1-\hat{\sigma})\bar{s}$. Solving the second row and comparing profits across cases then gives the equilibrium price, depending on the accountability level required for certification:

$$p^* = \begin{cases} u - (1-\hat{\sigma})\bar{s} & \text{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}, \\ \frac{u}{2} & \text{if } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \end{cases} \tag{A.6}$$

where the first row of (A.6) leads to a covered market, while some users do not buy cloud services in the second case. Substituting (A.6) in (A.5) yields the provider's equilibrium profits.

$$\pi^* = \begin{cases} u - (1-\hat{\sigma})\bar{s} - c(\sigma) - F & \text{if } \sigma = \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}, \\ \frac{u^2}{4(1-\hat{\sigma})\bar{s}} - c(\sigma) - F & \text{if } \sigma = \hat{\sigma} < 1 - \frac{u}{2\bar{s}}. \end{cases} \tag{A.7}$$

Recall that without certification the provider expects profits of $\pi(p^*) = \frac{u^2}{4\bar{s}}$. Consequently, she prefers to seek certification and to set accountability $\sigma = \hat{\sigma}$ over not seeking certification and setting $\sigma = 0$ if, and only if:

$$u - (1-\hat{\sigma})\bar{s} - c(\sigma) - F \geq \frac{u^2}{4\bar{s}} \quad \text{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}, \tag{A.8}$$

$$\frac{u^2}{4(1-\hat{\sigma})\bar{s}} - c(\sigma) - F \geq \frac{u^2}{4\bar{s}} \quad \text{if } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}. \tag{A.9}$$

Rearranging (A.8) and (A.9), certification is attractive for the provider if, and only if:

$$\hat{\sigma} \geq \frac{2(\bar{s}-u)^2}{4\bar{s}^2} + \frac{c(\sigma)+F}{\bar{s}} \quad \text{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}, \tag{A.10}$$

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

Both conditions in (A.10) state lower bounds on $\hat{\sigma}$. The binding (higher) condition is:

$$\hat{\sigma} \geq \begin{cases} 1 - \frac{u}{2\bar{s}} & \text{if } u \in (0, \frac{1}{2}(\bar{s} - \sqrt{5\bar{s}^2 - 8\bar{s}(c(\sigma) + F)})) \\ \frac{(\bar{s}-u)^2}{2\bar{s}^2} + \frac{c(\sigma)+F}{\bar{s}} & \text{if } u \in [\frac{1}{2}(\bar{s} - \sqrt{5\bar{s}^2 - 8\bar{s}(c(\sigma) + F}), 2\bar{s})) \end{cases} \quad \text{and } c(\sigma) + F < \frac{5}{8}\bar{s} \text{(A.12)}$$

In turn, equation (A.11) 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:

$$c(\sigma) + F < \frac{u}{2} + \frac{u^2}{4\bar{s}} \quad \text{for } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}. \quad \textit{Q.E.D.} \tag{A.13}$$

A.6 Proof of Lemma 6

Because the provider makes positive profits without certification—see Lemma 1—if the incentive compatibility constraints, (13) and (14), hold, the participation constraint also holds (and hence does not have to be considered explicitly any further). To find the profit-maximizing level of $\hat{\sigma}$ for the provider we solve (16) and find that it has a local maximum at:

$$\hat{\sigma} = \begin{cases} 1 - \frac{u}{2\sqrt{\bar{s}c'(\sigma)}} & \text{if } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \\ \sigma | c'(\sigma) = \bar{s} & \text{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}. \end{cases} \tag{A.14}$$

where $c'(\sigma)$ denotes the provider's marginal cost of accountability. The first row of (A.14) 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 (16) depends on the shape of $c(\sigma)$. This completely characterizes the solution to the providers' representatives' optimization problem.

By contrast, user representatives on the board maximize expected consumer surplus:

$$CS = \begin{cases} \int_0^{q(x)} (u - (1 - \hat{\sigma})s_i - p(x))ds_i & \text{if } x = 1, \\ \int_0^{q(x)} (u - s_i - p(x))ds_i & \text{if } x = 0. \end{cases} \tag{A.15}$$

Substituting $q(x = 1) = \overline{s_1}/\bar{s}$, $q(x = 0) = \overline{s_0}/\bar{s}$, and p^* from (A.6) we obtain:

$$CS = \begin{cases} \frac{1}{2}(2\bar{s} - 1)(1 - \hat{\sigma}) & \text{if } \sigma = \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}}, \\ \frac{(2\bar{s}-1)u^2}{8\bar{s}^2(1-\hat{\sigma})} & \text{if } \sigma = \hat{\sigma} < 1 - \frac{u}{2\bar{s}}, \\ \frac{(2\bar{s}-1)u^2}{8\bar{s}^2} & \text{if } \sigma = 0. \end{cases} \tag{A.16}$$

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

$$\frac{1}{2}(2\bar{s} - 1)(1 - \hat{\sigma}) > \frac{(2\bar{s} - 1)u^2}{8\bar{s}^2}, \quad \text{if } \hat{\sigma} \geq 1 - \frac{u}{2\bar{s}} \tag{A.17}$$

$$\frac{(2\bar{s} - 1)u^2}{8\bar{s}^2(1 - \hat{\sigma})} > \frac{(2\bar{s} - 1)u^2}{8\bar{s}^2}, \quad \text{if } \hat{\sigma} < 1 - \frac{u}{2\bar{s}}. \tag{A.18}$$

Inequality (A.17) 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 $\hat{\sigma}$, $1 - \frac{u}{2\bar{s}}$. Inequality (A.18) holds for all $0 < \hat{\sigma} < 1 - \frac{u}{2\bar{s}}$. The difference between the LHS and RHS of (A.18) is growing in $\hat{\sigma}$. Hence, in this case the user representatives prefer the highest supported level of $\hat{\sigma}$, for which $\hat{\sigma} < 1 - \frac{u}{2\bar{s}}$. *Q.E.D.*

References

Acquisti A, Varian HR (2005) Conditioning Prices on Purchase History. *Marketing Science* 24(3): 367-81.

Akerlof G (1970) The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics* 84(3): 488-500.

Aperjis C, Johari R (2010) Optimal windows for aggregating ratings in electronic marketplaces. *Management Science* 56(5): 864-880.

August T, Niculescu MF, Shin H (2014) Cloud Implications on Software Network Structure and Security Risks. *Information Systems Research* 25(3): 489-510.

Bakos Y, Dellarocas C (2011) Cooperation Without Enforcement? A Comparative Analysis of Litigation and Online Reputation as Quality Assurance Mechanisms. *Management Science* 57(11): 1944-1962.

Buehler B, Schuett F (2014) Certification and minimum quality standards when consumers are uninformed. *European Economic Review* 70: 493-511.

Cabral LMB, Hortacsu A (2010) The dynamics of seller reputation: Theory and evidence from eBay. *Journal of Industrial Economics* 58(1): 54-78.

Calzolari G, Pavan A (2006) On the Optimality of Privacy in Sequential Contracting. *Journal of Economic Theory* 130(1): 168-204.

Campbell J, Goldfarb A, Tucker C (2015) Privacy Regulation and Market Structure. *Journal of Economics & Management Strategy* 24(1): 47-73.

Casadesus-Masanell R, Hervas-Drane A (2015) Competing with Privacy. *Management Science* 61(1): 229-246.

Catteddu D, Hogben G (2009) An SME perspective on Cloud Computing. ENISA survey.

Chellappa RK, Sin RG (2005) Personalization versus Privacy: An Empirical Examination of the Online Consumer's Dilemma. *Information Technology and Management* 6: 181-202.

Cisco (2018) *Cisco Global Cloud Index: Forecast and Methodology, 2016–2021.* `https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.html.`

Corniere A de, De Nijs R (2016) Online Advertising and Privacy. *RAND Journal of Economics,* 47(1): 48-72.

Dellarocas C (2003) The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science* 49(10): 1407-1424.

Dellarocas C, Wood CA (2008) The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. *Management Science* 54(3): 460-476.

Dengler S, Prüfer J (2018) Consumers' Privacy Choices in the Era of Big Data. Tilburg University, TILEC Discussion Paper No. 2018-014.

Dixit AK (2003) Trade Expansion and Contract Enforcement. *Journal of Political Economy* 111(6): 1293-1317.

Dixit AK (2009) Governance Institutions and Economic Activity. *American Economic Review* 99(1): 5-24.

Edelman B (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 F (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 L, Prüfer J (2018) Faithful Strategies: How Religion Shapes Nonprofit Management. *Management Science*, forthcoming.

Forbes SJ, Lederman M (2013) Contract Form and Technology Adoption in a Network Industry. *Journal of Law, Economics, and Organization*, 29(2): 385-413.

Ganuza JJ, Gomez F, Robles M (2016) Product Liability versus Reputation. *Journal of Law, Economics, and Organization*, 32(2): 213-241.

Gibbons R, Henderson R (2012) Relational Contracts and Organizational Capabilities. *Organization Science* 23(5): 1350-1364.

Goldfarb A, Tucker C (2012) Shifts in Privacy Concerns. *American Economic Review: Papers and Proceedings* 102: 349-353.

Greif A, Milgrom PR, Weingast BR (1994) Coordination, Commitment, and Enforcement: The Case of the Merchant Guild. *Journal of Political Economy* 102: 745-776.

Haeberlen A (2010) A Case for the Accountable Cloud. *ACM SIGOPS Operating Systems Review* 44(2): 52-57.

Larrain M, Prüfer J (2015) Trade Associations, Lobbying, and Endogenous Institutions. *Journal of Legal Analysis* 7(2): 467-516.

McAfee (2017) Building Trust in a Cloudy Sky. Company report.

Masten SE, Prüfer J (2014) On the Evolution of Collective Enforcement Institutions: Communities and Courts. *Journal of Legal Studies* 43(2): 359-400.

Mathis J, McAndrews J, Rochet J-C (2009) Rating the raters: Are reputation concerns powerful enough to discipline rating agencies? *Journal of Monetary Economics* 56: 657-674.

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.

Pearson S (2011) Toward Accountability in the Cloud. *IEEE Internet Computing* 15(4): 64-69.

Persico N (2015) The Political Economy of Occupational Licensing Associations. *Journal of Law, Economics, and Organization*, 31(2): 213-241.

Peyrache E, Quesada L (2011) Intermediaries, Credibility and Incentives to Collude. *Journal of Economics & Management Strategy* 20(4): 1099-1133.

Probst CW, Sasse MA, Pieters W, Dimkov T, Luysterborg E, Arnaud M (2012) Privacy Penetration Testing: How to Establish Trust in Your Cloud Provider. *European Data Protection: In Good Health?*: 251-265.

Prüfer J (2013) How to Govern the Cloud? *IEEE CloudCom 2013* DOI 10.1109/CloudCom.2013.100: 33-38.

Prüfer J (2016) Business Associations and Private Ordering. *Journal of Law, Economics, and Organization*, 32(2): 306-358.

Rahman D (2012) But Who Will Monitor the Monitor? *American Economic Review* 102(6): 2767-2797.

Reed C (2013) Governance in Cloud Computing. Queen Mary University of London, Legal Studies Research Paper 157/2013.

Shaked A, Sutton J (1982) Relaxing Price Competition Through Product Differentiation. *Review of Economic Studies* 49: 3-13.

Shapiro C, Stiglitz JE (1984) Equilibrium Unemployment as a Worker Discipline Device. *American Economic Review* 74(3): 433-444.

Soghoian C (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): 360-424.

Solove DJ (2007) 'I've Got Nothing to Hide' and Other Misunderstandings of Privacy. *San Diego Law Review* 44: 745-772.

Spence M (1973) Job Market Signaling. *Quarterly Journal of Economics* 87(3): 355-374.

Stahl K, Strausz R (2011) Who should pay for certification? ZEW Discussion Paper 11-054.

Strausz R (2005) Honest certification and the threat of capture. *International Journal of Industrial Organization* 23: 45-62.

Taylor CR (2004) Consumer Privacy and the Market for Customer Information. *RAND Journal of Economics* 35(4): 631-50.

Tsai J, Egelman S, Cranor L, Acquisti A (2011) The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research* 22: 254-268.

Williamson OE (1999) Public and Private Bureaucracies: A Transaction Cost Perspective. *Journal of Law, Economics, & Organization* 15(1): 306-342.

Williamson OE (2002) The Lens of Contract: Private Ordering. *American Economic Review P&P* 92(2): 438-443.