UNIVERSITÀ DEGLI STUDI DI SIENA

QUADERNI DEL DIPARTIMENTO
DI ECONOMIA POLITICA E STATISTICA

Fabio Landini

The Evolution of Control in the Digital Economy

n. 655 – Ottobre 2012

Abstract - Control is becoming increasingly frequent in cyberspace, to an extent that puts into question the latter's traditional openness. In order to investigate the origins and effects of such change the paper formally model the historical evolution of digital control. In the model, the economy-wide features of the digital space emerge as a result of endogenous differences in culture (users' preferences including motivation) and technology (platform designs). The model shows that: a) in the long-run there exist two stable cultural-technological equilibria in the digital economy: one with intrinsically motivated users and low control; and the other with purely extrinsically motivated users and high control; b) under a closed economy - i.e. before the opening of the network to commerce, the initial emergence of a low-control-intrinsic-motivation equilibrium can be explained by the specific set of norms and values that formed the early culture of the networked environment; and c) the opening of the network to commerce can indeed cause a transition to a high-control-extrinsic-motivation equilibrium, even if the latter is Pareto inferior. Although it is too early to say whether such a transition is actually taking place, these results call for a great deal of attention in evaluating policy proposals on Internet regulation.

JEL Classification Numbers: C73; D02; K00; L23.

Keywords: Internet control, Internet regulation, motivation, on-line law enforcement, technology, endogenous preferences, evolutionary games

Fabio Landini, Department of Economics, University of Siena, email: landini4@unisi.it

1 Introduction

Control - i.e. the ability of those in power to direct and constraint the actions of others, is a crucial feature of many social institutions. Influential scholars from different disciplines - such as Marx (1970); Parsons (1963); Weber (1978) - have indeed dedicated a great deal of intellectual effort in understanding both its role and consequences. Significant emphasis, in particular, has been placed on the different modalities in which control could be exercised, with the distinction between psychological (Deci and Ryan, 1985), legal (Simon, 1951) and technological factors (Bowles, 1985).

Despite some notable early exceptions (e.g. Lessig, 1996; Reidenberg, 1998; Mitchell, 1995), the digital economy - i.e. the set of economic and social transactions that take place over the Internet, has been for long time considered as a place that is relatively exempt from strong forms of control. The reasons have been often associated with both the high cost of on-line rules enforcement (Johnson and Post, 1996; Elkin-Loren and Salzberger, 2000) and the advantage that loose forms of control may sometime have in sustaining innovation (Benkler, 2002a; von Hippel, 2005). It is not by chance the some of the most successful platforms in the last two decades have all included the giving up of control on some (or most) of their users' actions as key component of their internal design.[1]

Recently, however, there has been several signs of a turnround in the implementation of digital control. At the national level, for instance, there has been frequently reported cases of public authorities increasing their actual control over the Internet, going from the implementation of surveillance system to improve on-line security[2], to the recent controversy on Wikileaks's shutdown (Benkler, 2012a). At the corporate level, similarly, companies such as Google, Facebook and Apple have been all repeatedly accused of privacy and/or free-speech law infringements in because of the tight control they exercise on their digital platforms (MacKinnon, 2012). At the policy level, finally, three of the most recent and important initiatives to reform the Internet governance, such as the Anti-Counterfeiting Trade Agreement (ACTA)[3], the Stop On-line Piracy Act (SOPA) and Preventing Real On-line Threats to Economic Creativity and

[1] The intentional absence of control over users' actions (e.g. in the provision of content) is a key feature of most sharing-based platforms such as Wikipedia, YouTube, Flickr as well as the communities of free software developers and peer-to-peer file sharing networks. Similarly, the lack of control plays an important role also in the decentralized mechanisms of relevance and accreditation that are implemented in on-line marketplaces such as Amazon and eBay. All these platforms can be generally considered as instances of what Benkler (2006) calls peer production. For a detailed discussion of the role that users' decisional autonomy plays in peer production see Benkler and Nissenbaum (2006).

[2] See Robert Booth, *Government plans increased email and social network surveillance*, The Guardian, April 1, 2012, available at: `http://www.guardian.co.uk/world/2012/apr/01/government-email-social-network-surveillance` (last time checked: April 24, 2012).

[3] For a detailed analysis of ACTA and related criticisms see McManis (2008).

Theft of Intellectual Property Act (PIPA) in the U.S.[4] and the Google-Verizon Proposal on network neutrality[5], all look at control over users' on-line actions as a critical component of the proposed legislative framework. Overall, what seems to be undeniable, is that at present the digital economy has become radically different from the open and anarchic place it was at its origin. The effective direction in which it will further evolve, however, is still an open question; one whose answer is likely to have a direct impact on the way in which we, as citizens, contribute to society.

In this paper, I propose a behavioral economic model to study the evolution of digital control. By *digital control* I mean any coerced limitation of users' action which is imposed and enforced by the mean of digital code. In this sense I follow Lessig (1999, 2006) in considering *code* as the main instrument of rules enforcement in the digital space.[6] According to this definition, the deletion of a user's account because it is has proven responsible for illegal or controversial activities (e.g. diffusion of viruses, spamming, copyrights infringement) is a form of digital control.[7] The automatic removal of copyrighted files (e.g. songs, e-books) once they have been copied for a fixed number of times is another.[8] The discretionary choice of a marketplace owner concerning the types of applications that can be uploaded in her platform is a third one.[9] Obviously, these forms of control can be introduced and enforced at different layers of the overall Internet architecture (Lessig, 1999). For the sake of simplicity, I will focus only on the forms of control that can be enforced at the content layer, i.e. those that directly affect the organization of information production.[10]

I present a dynamic model where a group of platform designers and of users interact to produce information. Production is governed by the designers' choice on the degree of control, where the latter presents both costs and benefits. A

[4]See Jonathan Weisman, *After an Online Firestorm, Congress Shelves Antipiracy Bills*, The New York Times, January 20, 2012, available at:`http://www.nytimes.com/2012/01/21/technology/senate-postpones-piracy-vote.html?_r=1` (last time checked: April 24, 2012).

[5]See Claire Cain Miller and Miguel Helft, *Web Plan From Google and Verizon Is Criticized*, The New York Times, August 9, 2010, available at:`http://www.nytimes.com/2010/08/10/technology/10net.html?_r=2&ref=technology` (last time checked: April 25, 2012). For more detail on the concept of "net neutrality" see Wu (2003a).

[6]Lessig (1999, 2006) coined the well-known catchphrase "Code is Law" capturing the idea that in the digital space software code - as opposed to law, market and social norms - becomes the most powerful regulators of all. This is due to two main factors: first, the weaknesses of traditional law as a tool of on-line regulation; and second, the specific features of code that are associated with its malleability and nearly perfect enforceability. Overall, it is the combination of these specific features of code that, according to Lessig, makes cyberspace an arena of (potentially) perfect control.

[7]Similar provisions are included in the terms of service of most digital platforms, see for instance art. 5.5 in Facebook's Terms of Service: "if you repeatedly infringe other people's intellectual property rights, we will disable your account when appropriate", available at `http://www.facebook.com/legal/terms` (last time checked: April 25, 2012).

[8]See Zittrain (2000) on the creation of so-called trusted systems.

[9]See MacKinnon (2012) on Apple's App censorship practices.

[10]For a similar approach see Benkler (2002b).

high degree of control limits the set of tasks that can be performed by users and thus reduces the probability that a noisy activity is undertaken, where by *noisy activity* it is meant a task that may cause designers to incur a positive cost (e.g. illegal redistribution of copyrighted material, diffusion of worms and viruses, submission of malfunctioning applications). At the same time, control affects the users' preferences (and the behavior they support) for contributing information, in that it crowds out intrinsic motivation (e.g. desire for self-expression, ethical values, social norms). When intrinsic motivation is an important determinant of individual behavior this may generate a trade-off in the use of control. Given this setting, I study the evolution of control through a two-designers/two-users dynamic model. Designers can choose between two types of design, characterized by either high or low control. I refer to such distinction as an technological difference. The relevant preference differences are instead captured by assuming that users are either intrinsically or purely extrinsically motivated. The digital space is initially modeled as a closed economy, with no access from the outside. Then, a positive rate of accesses is allowed. This shift is aimed at capturing the opening of the network to commercial uses occurred in 1995. By studying the long-run equilibrium distributions of preferences and designs under these two different settings, I can make sense the overall evolution of digital control from the origin of cyberspace up to the present days. On this basis, I can also make some predictions on future trends.

This approach is characterized by two main novelties. The first one is that it takes into account the long-run effect of control on the distribution of on-line users' preferences. Most of the literature on digital control, especially on the legal side, has tended to neglect this effect. The attention, on the contrary, has been placed on the direct costs and benefits of control for information production. On the side of benefits, for instance, Zittrain (2000, 2008) extensively discusses the role that digital control can play in both reducing the degree of information noise at the content layer and improving the security of on-line transactions. For what concerns the costs, Lessig (2006) and Benkler (2002b) widely discuss the economic and political costs associated with a limitation of users' desires to access and redistribute data. On this respect, the present paper introduces an additional effect of control which consists of its influence on the type of culture (i.e. preference distribution) that characterizes on-line participation. As we will see, under certain conditions, this effect may also have normative implications.

The second novelty of this approach is that, instead of treating control and motivations as exogenous or determined by an economy-level institutional bargain, it exploits evolutionary game theory to model the interacting dynamics of both as the result of decentralized non-cooperative interactions among agents.

In doing so the paper adopts a methodological approach which is similar to Bowles et al. (2003), Naidu et al. (2010), Belloc and Bowles (2011), Bisin and Verdier (2001) and Landini (2012). None of this paper, however, studies the evolution of digital control.

On the basis of the model, I derive three main results. First, in the long-run there exist two stable cultural-technological equilibria in the digital economy: one in which all users are intrinsically motivated and designers exercise low control; and the other in which all users are purely extrinsically motivated and designers exercise high control. In this sense, the model reflects Lessig's view on the existence of two extremely different social spaces the digital economy could eventually evolve into. The emergence of one of these two spaces depends on several factors, among which the actual costs and effectiveness of control.

Second, during the early days of the digital economy (i.e. period 1969-1995) the closure of the network to commercial uses favored both the emergence and persistence of a low-control-intrinsic-motivation equilibrium. In that period, in fact, most of the network users looked at the emerging public Internet more as an instrument to enable free and open communication rather than as a tool for running businesses, and were thus characterized by fairly strong intrinsic motivation. This, combined with the relatively high costs of digital control, favored the initial emergence of a cultural-technological equilibrium dominated by low-control designs. At the same time, the ban to exploit the network for commercial purposes, transformed the digital space in a sort of closed system, characterized by relatively few accesses from the outside. This closure imposed a limit on the possibility that some forms of idiosyncratic shock could induce a transition to a different type of equilibrium, thus sustaining the persistence of the low control status-quo.

Third, the opening of the network to commercial uses can indeed cause the transition to a different type of cultural-technological equilibrium. In particular, by allowing for a positive rate of exogenous variation, the long-run effect of such opening may be to displace the low-control-intrinsic-motivation equilibrium in favor of the high-control-extrinsic-motivation one. Quite interestingly, I find that such displacement may occur even if the high-control-extrinsic-motivation equilibrium is Pareto inferior. Although it is too early to say whether such a transition is actually taking place, this result calls for a great deal of attention in evaluating the role and pace of government intervention in the digital space.

The paper is related with two main streams of literature. The first one is the law and economics literature on rule-making in cyberspace, which includes the seminal contributions by Johnson and Post (1996), Lessig (1996, 1999), Post (1995), Mitchell (1995) and Reidenberg (1998), as well as more recent works by Wu (2003b), Zittrain (2003, 2006, 2008), Strahilevitz (2003), Goldsmith and

Wu (2006), and Deibert et al. (2010). While these contributions are mainly concerned with the governance of the overall Internet architecture, this paper exploits some insights from these works - namely the idea of code as an efficient enforcement device - to study the organization of information production at the content layer. It does so by adopting a bottom-up, dynamic and emergence-based approach to the analysis of institutional change in the digital space. In this way the paper directly addresses the call by Elkin-Loren and Salzberger (2000) for new approaches to the study of economic institutions in cyberspace.

In its behavioral assumptions the paper is also related to the social psychology and behavioral economics literature on intrinsic and extrinsic motivation (Deci and Ryan, 1985; Frey, 1997; Frey and Jegen, 2001). Such literature has provided solid empirical and experimental evidence supporting both the role of intrinsic motivation and the existence of motivational crowding out as a result of exogenous incentives - for a recent survey of the empirical results see Bowles and Polania-Reyes (2012). Although the interplay between different types of motivation has been already considered in formal economic modeling (Benabou and Tirole, 2003), less emphasis has been placed on the interaction between control and motivation in the digital space. This aspect, on the contrary, is at the core of the present paper.

The paper is organized as follows. Section 2 describes the setting of the model and defines the main assumptions. Section 3 introduces the model's dynamics and finds the equilibrium conditions in a closed economy (i.e. for the period 1969-1995). Section 4 adds the possibility of external entrance and studies stochastic stability. Section 5 discusses the major policy implications. Section 6, finally, concludes.

2 Code, control and motivation

A digital economy is populated by n_d (> 0) platform designers (d) and n_u (> 0) users (u), with $n_d < n_u$. Each designer owns a platform and repeatedly interact with users to produce information, the single interaction being a random user-designer match in which a generic user i performs action a_i being offered a given design. Designs differ according to the degree of digital control they support. The payoffs of users and designers are modeled as follows.

A digital platform is represented as a set of tasks. For any given platform type (e.g. social networks, wikis, peer-to-peer services, apps store), I call T the set of all feasible tasks that are enabled by the features of the available technology (e.g. upload and/or download of files, interaction with other users, provision of particular on-line services). Given T, and the specific interests of the platform designer, I assume the existence of a non-empty set of noisy tasks $T_k \subseteq T$ such that if any $t_k \in T_k$ is performed by a user, the platform designer incurs a cost k (> 0). Such tasks may include, for instance, the illegal redistribution of copyrighted material, where k corresponds to the state-enforced sanction on the platform owners that are deemed to facilitate copyright law infringement. Another example consists of undesired users' messaging, where k corresponds to the (designer's expected) reduction in the number of accesses to the platform caused by an increased noisiness of the information environment. In order to avoid this cost, the platform designer may choose to define a set of censored tasks $T_{ce} \subseteq T_k$, such that any $t_{ce} \in T_{ce}$ cannot be performed by any user. Such censoring is enforced by the mean of digital code, in the sense that all $t_{ce} \in T_{ce}$ are simply made "not available" to users (see Lessig, 2006). On this basis, the degree of digital control supported by the platform design can be represented by the ratio $t = |T_{ce}|/|T_k| \in [0, 1]$ where $|T_{ce}|$ and $|T_k|$ are the cardinality of sets T_{ce} and T_k respectively.

The degree of digital control t plays two main roles in the model. On one hand it affects the probability that a noisy task $t_k \in T_k$ is performed. In particular, the higher t, the lower such probability. On the other, it affects the user's motivation for performing action a_i, where the latter is defined as the effort exercised in performing some (or all) of the available tasks. Tasks availability obviously depends on the degree of censoring, and can be represented by set $T_{av} = T \setminus T_{ce}$. When there is no censoring (i.e. $T_{ce} = \oslash$ and $t = 0$), the set of available tasks coincides with the set of feasible tasks, so that $T_{av} = T$. On the contrary, when censoring is maximum (i.e. $T_{ce} = T_k$ and $t = 1$), the set of available tasks contains only "non-noisy" tasks and $T_{av} = T \setminus T_k$.

Given these definitions, I model individual motivation as follows. Each user i can obtain two main types of reward from performing action a_i. The first type consists of the extrinsic benefits associated with *participation in* the platform, and (depending on the type of platform) includes things such as (present or delayed) monetary rewards, increased reputation and access to fast communication tools. The second type of reward coincides with the non-monetary and intrinsic benefits associated with the *contribution to* the platform, and (again depending on the type of platform) includes the pure pleasure of information sharing, the desire for self-expression, and the utility derived from cooperating with others.

In line with the results reported in several studies coming from both social psychology and behavioral economics, I assume that the degree of digital control is not neutral with respect to the nature of individual motivation. In particular, I assume that an increase in t has two main effects. First, it crowds out intrinsic motivation, i.e. for any given level of extrinsic motives, it reduces the total marginal benefit that a user derives from performing a_i.[11] In the off-line world, this effect has been found in a large number of natural environments (Gagne and Deci, 2005) and experimental settings (Irlenbusch and Ruchala, 2008), and there are reasons to believe that it holds also on-line. Second, I assume that an increase in t generates a net positive disutility for the users with intrinsic motives, whose value is independent of a_i. On this respect, some supporting evidence comes from the growing experimental economics literature on control-aversion (Falk and Kosfeld, 2006; Fehr et al., 2010; Charness et al., 2011).

Formally, I call ϕ and $\lambda_i(t)$ the extrinsic and intrinsic marginal benefit of a_i. In line with the crowding-out hypothesis, I assume $\lambda_i(0) = \lambda_i > 0$ and $\lambda_i' < 0$. While ϕ is user-generic, I consider $\lambda_i(t)$ to be user-specific. This captures the idea that while most individuals are motivated by some forms of extrinsic reward, only some of them exhibit also intrinsic motivation. As a behavioral model, this is a fair compromise between the standard economic view of self-interested and purely extrinsically motivated agents, and the behavioral and psychological approach based on a more complex mix of motivations. With respect to the costs, I call $c(a_i)$ the opportunity cost of the time spend in performing a_i and $\mu_i(t)$ the psychological cost associated with control aversion, where $c' > 0$, $c'' > 0$ and $\mu_i' > 0$. On this basis, following Bowles and Hwang (2008), I assume i's (risk-neutral) utility function to be additive in motivations, and I write:

$$U_i = [\phi + \lambda_i(t)]a_i - c(a_i) - \mu(t) \tag{1}$$

From the maximization of Eq. 1 with respect to a_i it follows that:

Remark 1 The optimal level of a_i is given by condition $c' = \phi + \lambda_i(t)$. Since $\lambda' < 0$, an increase in t reduces *ceteris paribus* the optimal level of a_i for intrinsically motivated users.

At an intuitive level, the relationship between motivation and control on one hand, and the level of a_i on the other can be thought in terms of number and typology of tasks being performed by users. Within set T_{av}, in fact, we can generally distinguish between two main typologies of tasks. On one hand there are tasks that can be directly exploited for extrinsic purposes. They include, for instance, the possibility to share information that advertise the user's last piece of work (e.g. song, book, or academic article depending on the platform). On the other, there exist "self-policing" tasks that tend to be associated only with the intrinsic pleasure of contributing to the platform's well-functioning. They

[11] This way of modeling motivational crowding out is generally called "marginal". An alternative is to assume "categorical" crowding out. On the distinction between marginal and categorical crowding out see Bowles and Polania-Reyes (2012).

include, in the order, the self-reporting of bugs and errors, the sanctioning of other users' misbehavior and the updating of missing information. On this basis, the positive relationship between the value of $\lambda_i(t)$ and the optimal level of a_i can be interpreted with the fact that, for any given level of t, intrinsically motivated users are willing to perform both typologies of tasks (and thus to exert high effort), whereas purely extrinsically motivated ones tend to undertake only the former. Starting from this condition, an increase in t tends to reduce the effort of intrinsically motivated users because by undermining the degree of self-commitment to the platform it decreases their willingness to undertake "self-policing" acts.

Given this behavioral model for users, I now consider the payoff of designers. Each designer owns a platform and earns a reward that is positively related to the amount of information produced by users. For any given user-designer match, I write the return of the designer as $q(a_i) = qa_i$, where $q > 0$ is the marginal contribution of i's action to the overall stock of information. The decision of each designer is then concerned with the degree of control t to be implemented. Depending on users' motivation, control presents both costs and benefits. On the side of costs, as reported in Remark 1, control reduces the level of a_i whenever the designer is matched with an intrinsically motivated user. I call the latter the motivational cost of control. Moreover, control forces the designer to spend both cognitive and digital resources (e.g. lines of code) to define the tasks to be included in T_{ce}. This can be termed the design cost of control and is represented by function $\delta(t)$, with $\delta' > 0$ and $\delta'' > 0$. On the side of benefits, on the contrary, control limits the number of tasks in T_k that are effectively available to users and thus reduces the probability of incurring cost k. Such probability obviously depends also on the willingness of users to undertake tasks included in set $T_k \setminus T_{ce}$, i.e. noisy tasks that are not censored by the designer. On this respect, given the association between intrinsic motivation and social/ethical norms, I will assume that such willingness is lower the more intrinsically motivated the user.

On this basis I write the expected reward of a generic (risk-neutral) designer d as follows:

$$\pi_d = qa_i - \delta(t) - \gamma(1-t)\eta(\lambda_i)k \tag{2}$$

where $\eta(\lambda_i) \in [0,1]\ \forall\ \lambda_i \geq 0$ with $\eta' < 0$ is the probability that i performs a task in T_k and $\gamma \in [0,1]$ represents the effectiveness of censoring in reducing the probability of incurring cost k. At an intuitive level, γ captures both the status of the control-enhancing technologies and the users' capabilities to hack the limitations imposed by designers, where the more effective the former and weaker the latter, the higher the value of γ.

Before going deeper into the analysis, I want now to point out three key assumptions that are related to the way in which users and designers interact in the economy. First, I assume that there are no strategic interactions among users. Every period, a user decides on her optimal level of a_i by looking only at the degree of t implemented by the designer she is matched with, and on the basis of the preferences described by Eq. (1). In this way I abstract from any form of free-riding problem that may arise when users simultaneously contribute to the same platform. Although this is a fairly strong assumption, it heavily simplifies the analysis and it allows me to focus on the motivational effect of control. Moreover, several studies have shown that the incentives to free-ride do not represent a big problem in most digital platforms, especially when the standard money maximizer behavioral model is expanded to include also intrinsic motivation (see Benkler, 2006). On this basis, I simply choose to take free-riding incentives out of the analysis.

Second, I assume that there is no assortment in the matching between users and designers. In other words, neither users nor designers can choose the type of partner to interact with. This is due to the fact that both intrinsic motivations and digital control are assumed to be unobservable ex-ante, and thus cannot be used to condition the matching dynamics. This is a standard assumption in most evolutionary game theoretic model.

Finally, I assume that while both motivation and control change over time, none is the result of instantaneous individual maximization. Rather, they are durable features of users and designers that evolve in a decentralized environment under the influence of long-run economy-wide payoff differences. Users, in particular, periodically update their motivation by best responding to the distribution of designs in the past. Similarly, designers occasionally update the degree of control by best responding to the past distribution of users' motivation. The main objective of the analysis is then to identify the long-run equilibrium distribution of both motivation and designs.

3 The closed economy, 1969-1995

3.1 Stage game

Given the setting described in Section 2, I now study the evolution of digital control under the assumption of no external access into the economy. In every period, the members of both populations of users and designers remain the same, and interactions evolve through their continuous re-matching. Such an assumption is aimed at capturing the status of the digital economy as of the period 1969-1995, when the prohibition to use the public Internet for commercial purposes brought a relative closure of cyberspace. Such an assumption will be later removed (see Section 4), when I study the opening of the network to commerce.

From the technical point of view, I also introduce three important simplifications that will be maintained throughout the model. First, I assume an explicit form for functions $\lambda(t)$, $\mu(t)$, $\delta(t)$, $c(a)$ and $\eta(\lambda)$. In particular, I assume $\lambda(t) = \lambda(1-t)$ with $\lambda > 0$, $\mu(t) = \mu t$ with $\mu > 0$, $\delta(t) = \delta t^2/2$, $c(a) = a^2/2$ and

$$\eta(\lambda) = \begin{cases} 1, & \text{if } \lambda = 0 \\ \\ \eta, & \text{if } \lambda > 0 \end{cases} \tag{3}$$

where $0 < \eta < 1$. This simplifies the analysis, without affecting the final results. Second, I assume that there exist only two available control technologies in the economy, namely a full-control ($t = 1$) and a no-control ($t = 0$) technology. Designers employing the full-control technology are called high control designers (H-type), whereas those employing the no-control technology are called low control designers (L-type). This distinction is obviously an oversimplification, which however makes the model analytically traceable. Finally, I assume that users are either purely extrinsically motivated (E-type), or both extrinsically and intrinsically motivated (I-type). Within the population of I-users I assume that no difference exists in terms of the degree of intrinsic motivation, so that there are only two behavioral types. Looking at the explicit functions defined above, such types can be defined by the pair (λ, μ), with I-user $\sim (\lambda, \mu)$ and E-user $\sim (0, 0)$. This, when combined with Eq. (1), gives us the following utility functions for I- and E-users respectively:

$$U_I = [\phi + \lambda(1-t)]a - \frac{a^2}{2} - \mu(t) \quad \textit{and} \quad U_E = \phi a - \frac{a^2}{2} \tag{4}$$

We can thus derive the following (proofs for all Lemmas and Propositions are reported in Appendix A):

Lemma 1. Call $a_{i,j}$ the best-response level of a for an i-type user when matched with a j-type designer. From Remark 1 and Eqs. (4) it follows that: $a_{I,L} = \phi + \lambda$ and $a_{I,H} = a_{E,H} = a_{E,L} = \phi$.

Given these simplifications, it is now possible to represent the single interaction taking place between a user u and a designer d in game theoretic form as follows. Let's interpret each behavioral and technological type as the strategy of a game, which is usually called the stage game. We thus have $\Sigma_u = \{I, E\}$ and $\Sigma_d = \{L, H\}$, where $\Sigma_i =$ (for $i = u, d$) is the strategy set of player i.

Designers ($\rightarrow$) Users($\downarrow$)	**L-type** ($t = 0$)	**H-type** ($t = 1$)
I-type $\sim (\lambda, \mu)$	$\frac{(\phi + \lambda)^2}{2}$, $q(\phi + \lambda) - \gamma\eta k$	$\frac{\phi^2}{2} - \mu$, $q\phi - \frac{\delta}{2}$
E-type $\sim (0, 0)$	$\frac{\phi^2}{2}$, $q\phi - \gamma k$	$\frac{\phi^2}{2}$, $q\phi - \frac{\delta}{2}$

Table 1: Stage game matrix of payoffs. Note: each cell of the matrix represents a different preference-design matching that may occur in the economy.

On this ground, a stage game of information production can be defined by the triplet $\Gamma = (I, \Sigma, \pi)$ where $I = \{u, d\}$ is the set of players, $\Sigma = \Sigma_u \times \Sigma_d$ is the set of strategy profiles and $\pi = \{\pi_u(\sigma), \pi_d(\sigma)\}$ for $\sigma \in \Sigma$ is the vector function of players' payoffs, where $\pi_u(\sigma)$ and $\pi_d(\sigma)$ are given by Eqs. (1) and (2) respectively. Table 1 reports a normal-form representation of Γ, taking into consideration that functions $\lambda(t)$, $\mu(t)$, $\delta(t)$, $c(a)$ and $\eta(\lambda)$ take the explicit form defined above (for the derivation of the payoffs in Table 1 see Appendix B). With respect to Γ, I thus introduce the following definitions:

Definition 1. *A preference-design matching in game Γ corresponds to a pure strategy profile $\sigma = \{\sigma_u, \sigma_d\} \in \times_{i \in I} \Sigma_i$, where $\sigma_u \in \Sigma_u$ and $\sigma_d \in \Sigma_d$ is the pure strategy adopted by player u and d respectively.*

Definition 2. *A preference-design matching $\sigma^* = \{\sigma_u^*, \sigma_d^*\}$ is a preference-design equilibrium, if the correspondent pure strategy profile is a Nash equilibrium of game Γ.*

Game Γ offers a mapping of all possible preference-design matchings that may occur in the economy, namely: $\{I, L\}$, $\{E, L\}$, $\{I, H\}$ and $\{E, H\}$. I now determine the conditions under which each of the four matchings is also an equilibrium.

Proposition 1. *Suppose $\lambda > 0$, $0 < \eta < 1$, $k > 0$ and $\mu > 0$. Then, $\exists$ two values $\underline{\delta} = 2(\gamma\eta k - q\lambda)$ and $\overline{\delta} = 2\gamma k$ (for $\overline{\delta} > \underline{\delta}$) s.t.: (i) if $\delta > \overline{\delta}$, then $\{I, L\}$ is the only preference-design equilibrium; (ii) if $\delta < \underline{\delta}$, then $\{E, H\}$ is the only preference-design equilibrium; and (iii) if $\underline{\delta} < \delta < \overline{\delta}$, then there exist two preference-design equilibria, namely $\{I, L\}$ and $\{E, H\}$.*

Corollary 1.1 *For any value of δ, $\{E, L\}$ and $\{I, H\}$ are never preference-design equilibria.*

Corollary 1.2 *For any value of δ, there always exist at least one preference-design equilibrium in game Γ.*

At the population level a preference-design equilibrium represents a cultural-technological convention, meaning that conforming to it is a mutual best response as long as virtually all members of each population (users and designers) expect virtually all members of the other to conform to it. According to Proposition 1, the number and types of conventions existing in the economy depend on the design cost of control δ. When the latter is greater than an upper threshold $\overline{\delta}$ (because for instance technology is costly to manipulate), $\{I, L\}$ is the only cultural-technological convention in the economy, and is thus likely to proliferate. On the contrary, when δ is smaller than a lower threshold $\underline{\delta}$ ($< \overline{\delta}$), the only cultural-technological convention is $\{E, H\}$. Quite interestingly, I find that when δ is intermediate between these two values, two cultural-technological conventions exist in the economy, namely $\{I, L\}$ and $\{E, H\}$. In this case, the convention that will emerge as the long-run cultural-technological equilibrium of the economy depends on the asymptotic stability properties of the two conventions. From the analytical point of view, this is clearly the most interesting case to study.

Before going into the details of asymptotic stability, it is interesting to characterize the efficiency properties of the two conventions. In this sense I find that, if anything, convention $\{I, L\}$ tends to exhibit an efficiency advantage with respect to $\{E, H\}$. In particular, I derive the following result:

Proposition 2. *If in the economy there exist only one cultural-technological convention, then the associated preference-design equilibrium is Pareto efficient. If in the economy there exist two cultural-technological conventions, then $\{I, L\}$ Pareto dominates $\{E, H\}$*

The intuition behind Proposition 2 is straightforward, and essentially relates to the combined effect of λ and δ. Since $\lambda > 0$, users are always better-off under $\{I, L\}$, because they can enjoy the additional utility that derives from being intrinsically motivated. This implies that $\{I, L\}$ will never be Pareto dominated. At the same time, the payoff condition of designers depends on the specific value of δ. When δ is sufficiently low (i.e. $\delta < \underline{\delta}$), designers are better-off under $\{E, H\}$, because control is cheap and allows one to avoid cost k. In the latter case $\{E, H\}$ is the only convention of the economy and it is also Pareto efficient. When δ is sufficiently high (i.e. $\delta > \underline{\delta}$), on the contrary, designers are better-off under $\{I, L\}$, because the cost of designing an architecture of control is less than compensated by the reduced likelihood of incurring cost k. This implies that for the range of values in which both $\{I, L\}$ and $\{E, H\}$ are cultural-technological conventions (i.e. for $\underline{\delta} < \delta > \overline{\delta}$), the former Pareto dominates the latter. As we will see later in the paper, the Pareto dominance of $\{I, L\}$ over $\{E, H\}$ turns out to be a critical feature of the model, which has interesting policy implications.

3.2 Dynamics

To provide a framework for studying asymptotic stability I now restrict the analysis to the space of parameter in which two conventions exist (i.e. I assume $\underline{\delta} \leq \delta \leq \overline{\delta}$) and introduce an explicit model of the dynamics of change. In particular, I model such dynamics as follows. In every time period $d\tau$ users and designers are randomly paired to play the stage game described above. Give their own type and the degree of control chosen by the designer, users choose their level of effort according to the best response functions reported in Lemma 1. Once production has taken place, designers and users earn the payoffs reported in Table 1. Let:

$$\omega_I^\tau = \frac{n_I^\tau}{n_I^\tau + n_E^\tau} \quad \omega_L^\tau = \frac{n_L^\tau}{n_H^\tau + n_L^\tau} \tag{5}$$

be the fractions of I-users and L-designers operating in the economy at any τ, where n_i^τ (for $i = I, E, L, H$) is the number of agents (users and designers) of type i in period τ. The pair $\{\omega_I^\tau, \omega_L^\tau\}$ represents the state of the economy, i.e. it gives the overall distribution of motivation and control. Assuming that the size of the economy is sufficiently large, ω_I^τ and ω_L^τ will also denote the probability with which users and designers are paired across types. On this basis, for any given value of ω_L^τ and taking into consideration the payoffs reported in Table 1, we can write the expected payoffs of I- and E-users at any τ as follows:

$$V_I^\tau(\omega_L^\tau) = \omega_L^\tau \left[\frac{(\phi + \lambda)^2}{2}\right] + (1 - \omega_L^\tau) \left[\frac{\phi^2}{2} - \mu\right] \tag{6}$$

$$V_E^\tau(\omega_L^\tau) = \omega_L^\tau \frac{\phi^2}{2} + (1 - \omega_L^\tau) \frac{\phi^2}{2} \tag{7}$$

Similarly, for any given value of ω_I^τ, the expected payoffs to L- and H-designers are respectively:

$$V_L^\tau(\omega_I^\tau) = \omega_I^\tau \left[q(\phi + \lambda) - \gamma\eta k\right] + (1 - \omega_I^\tau) \left[q\phi - \gamma k\right] \tag{8}$$

$$V_H^\tau(\omega_I^\tau) = \omega_I^\tau \left[q\phi - \frac{\delta}{2}\right] + (1 - \omega_I^\tau) \left[q\phi - \frac{\delta}{2}\right] \tag{9}$$

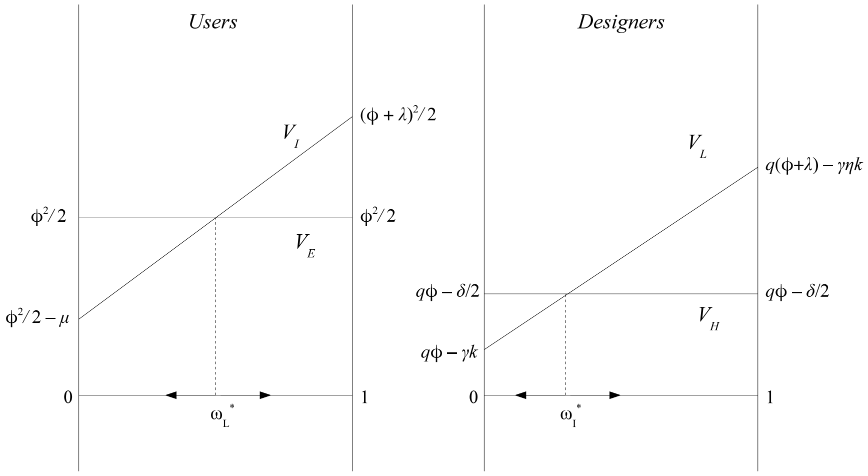

Figure 1: Expected payoffs to I- and E-users and L- and H-designers. Note: ω_I^τ is the fractions of users who are intrinsically motivated and ω_L^τ the fraction of designers who exercise low code-based control at time τ. The vertical intercepts are from Table 1.

These expected payoff functions are illustrated in Figure 1.

To model the co-evolution of motivations and control, suppose that both users and designers update the preferences and the designs (respectively) by best responding to the distribution of types in the previous period. In particular, suppose the updating process works as follows. In any time period $d\tau$ both users and designers are exposed to a cultural or technological model randomly selected from their subpopulation. For instance, a designers, named A, has the opportunity to observe the degree of control exercised by another designers, named B, and to know her expected payoff with a probability $\alpha d\tau$. If B is the same type as A, A does not update. But if B is of a different type, A compares the two payoffs and, if B has a greater payoff, switches to B's type with a probability equal to β (> 0) times the payoff difference, retaining her own type otherwise. The same procedure takes place among users. Specifically, writing the probability that an agent (user and designer) of type i switches to type j at time τ as p_{ij}^τ we have:

$$p_{ij}^\tau = \begin{cases} \beta\left(V_j^\tau - V_i^\tau\right), & \text{if } V_j^\tau > V_i^\tau \\ \\ 0, & \text{if } V_j^\tau \leq V_i^\tau \end{cases} \tag{10}$$

for $i, j = I, E$ and $i \neq j$ in the case of users and $i, j = L, H$ and $i \neq j$ in the case of designers. On this basis the expected fractions of I-users in period $\tau + d\tau$ is given by:

$$\omega_I^{\tau+d\tau} = \omega_I^\tau - \omega_I^\tau(1-\omega_I^\tau)\alpha d\tau\sigma_E\beta(V_E^\tau - V_I^\tau) + (1-\omega_I^\tau)\omega_I^\tau\alpha d\tau\sigma_I\beta(V_I^\tau - V_E^\tau) \quad (11)$$

where σ_E and σ_I are two binary functions such that $\sigma_E = 1$ if $V_E^\tau > V_I^\tau$ and is zero otherwise, $\sigma_I = 1$ if $V_I^\tau \geq V_E^\tau$ and is zero otherwise, and $\sigma_E + \sigma_I = 1$. Eq. (11) reads as follows: the expected fraction of I-users at $\tau + d\tau$ is given by the fraction of I-users at τ (first term), minus the fraction of I-users who are paired with an E-user and switch their type (second term), plus the fraction of E-users who are paired with an I-user and switch their type (third term). Similarly, the expected fractions of L-designers in period $\tau + d\tau$ is given by:

$$\omega_L^{\tau+d\tau} = \omega_L^\tau - \omega_L^\tau(1-\omega_L^\tau)\alpha d\tau\sigma_H\beta(V_H^\tau - V_L^\tau) + (1-\omega_L^\tau)\omega_L^\tau\alpha d\tau\sigma_L\beta(V_L^\tau - V_H^\tau) \quad (12)$$

where $\sigma_H = 1$ if $V_H^\tau > V_L^\tau$ and is zero otherwise, $\sigma_L = 1$ if $V_L^\tau \geq V_H^\tau$ and is zero otherwise, and $\sigma_H + \sigma_L = 1$. Subtracting ω_I^τ and ω_L^τ from both sides of Eqs. (11) and (12) respectively, dividing both equations by $d\tau$, and taking the limit as $d\tau \to 0$, we get:

$$\dot{\omega}_I^\tau = \omega_I^\tau(1 - \omega_I^\tau)(V_I^\tau(\omega_L^\tau) - V_E^\tau(\omega_L^\tau)) \quad (13)$$

$$\dot{\omega}_L^\tau = \omega_L^\tau(1 - \omega_L^\tau)(V_L^\tau(\omega_I^\tau) - V_H^\tau(\omega_I^\tau)) \quad (14)$$

where, for the sake of simplicity, I have assumed $\alpha\beta = 1$. Eqs. (13) and (14) represent a system of differential equations which describes how the distribution of types $\{\omega_I^\tau, \omega_L^\tau\}$ evolve over time. Given this dynamics, we are mainly interested in the stationary states of the economy, namely the states for which $\dot{\omega}_I^\tau = 0$ and $\dot{\omega}_L^\tau = 0$. Such states represents fixed-points of the dynamical system, and cultural-technological equilibria of the economy.

Proposition 3. *The dynamical system composed of Eqs.* (13) *and* (14) *is characterized by five cultural-technological equilibria:* $\{0,0\}$*,* $\{0,1\}$*,* $\{1,0\}$*,* $\{1,1\}$ *and* $\{\omega_I^*, \omega_L^*\}$*, where*

$$\omega_I^* = \frac{2\gamma k - \delta}{2\left[q\lambda + \gamma k(1-\eta)\right]} \qquad \omega_L^* = \frac{2\mu}{\lambda(\lambda + 2\phi) + 2\mu} \quad (15)$$

Out of these five equilibria, only two are asymptotically stable, namely $\{0,0\}$ *and* $\{1,1\}$*; equilibrium* $\{\omega_I^*, \omega_L^*\}$ *is a saddle, whereas equilibria* $\{0,1\}$ *and* $\{1,0\}$ *are unstable*

Because $\underline{\delta} \leq \delta \leq \overline{\delta}$, both ω_I^* and ω_L^* are included in the close interval $[0,1]$.

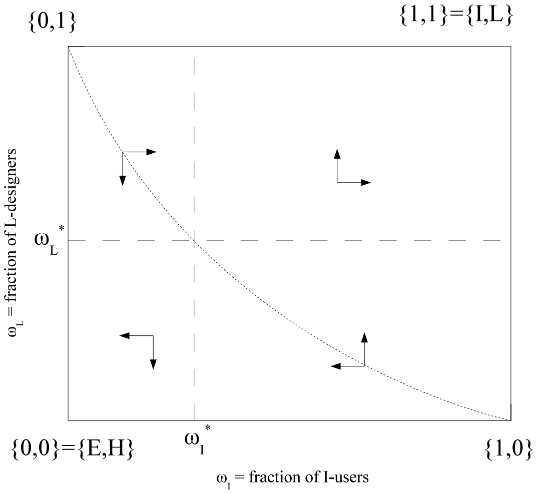

Figure 2: Asymptotically stable states and out-of-equilibrium dynamics. Note: the arrows represent the disequilibrium adjustment in the number of I-users (horizontal movements) and L-designers (vertical movements).

The vector field in Figure 2 offers a graphical representation of the dynamical system composed of Eqs. (13) and (14), and of the content of Proposition 3. The arrows indicate the out-of-equilibrium adjustment. For states $\omega_I^\tau < \omega_I^*$ and $\omega_L^\tau < \omega_L^*$ (i.e. in the southwest region of Figure 2), both $\dot{\omega}_I^\tau$ and $\dot{\omega}_L^\tau$ are negative and the economy will move to $\{0,0\}$. This state corresponds to a cultural-technological equilibrium in which H-type designers interact with E-type users; I will call the latter Equilibrium 0 (E_0). Analogous reasoning holds for the northeast region of Figure 2, where the economy converges to $\{1,1\}$. In this case the stable state corresponds to a cultural-technological equilibrium in which L-type designers interact with I-type users; I will call the latter Equilibrium 1 (E_1). In the remaining regions of the state space, namely northwest and southeast, we may identify a locus of states (dashed downward-sloping line) for which the system will transit to the interior equilibrium $\{\omega_I^*, \omega_L^*\}$, with states below that locus transiting to E_0, and above the locus to E_1. State $\{\omega_I^*, \omega_L^*\}$ is stationary, but is a saddle: small movement away from it are not self-correcting. Two additional unstable stationary states are $\{1,0\}$ and $\{0,1\}$, but are of no interest. All the area below the dashed downward-sloping line represents instead the basin of attraction of E_0, and all the area above it the one of E_1. These two corner solutions are thus the absorbing states of the dynamic process. If the economy is ever at either of these states, it will never leave.

The dynamics represented in Figure 2 suggests that, overtime, the economy is likely to converge to one of two very different equilibria. In one of them, namely E_0, a homogeneous population of extrinsically motivated users interact overtime with designers employing high control technologies. In the other, namely E_1, a population dominated by intrinsically motivated users interact with designers exercising low control. According to Proposition 2, the convergence to one equilibrium as opposed to the other does indeed have implications in terms of overall efficiency, in that E_1 is Pareto dominant over E_0. The extend to which one of these two equilibria will actually be the cultural-technological equilibrium of the economy depends on two interrelated factors. First of all, for any given size of the basins of attraction, the emergence of E_1 as opposed to E_0 (and *viceversa*) is more likely, the more probable the initial distribution of types in the economy to fall within E_1's (or E_0's in the opposite case) basin of attraction. This implies that, in this closed setting, there exist path dependency in the way in which the economy evolves. Secondly, for any given initial distribution of types, the emergence of one of the two absorbing states as the final resting point of the dynamics depends on the size of its basin of attraction. In particular, the greater the basin of attraction of one state relative to the other, the more likely such state to become the cultural-technological equilibrium of the economy. On this respect, it is important to notice that:

Remark 2. $\partial\omega_I^*/\partial\delta < 0$ and $\partial\omega_I^*/\partial\gamma > 0$ imply that, for any initial distribution of types, the emergence of E_1 as the cultural-technological equilibrium of the economy is more likely, the greater the design cost of control δ and the less effective the technologies of censoring, i.e. the lower γ.

3.3 Discussion

Looking at the features of the digital economy at the end of the 1969-1995 period, it is easy to see which type of cultural-technological equilibrium the non-commercial network eventually evolved into. Several authors indeed agree in considering both the relatively low degree of control and the widespread diffusion of intrinsic motives and social norms as two peculiar features of the networked environment as of the mid-1990s (Zittrain, 2008; Lessig, 2006; Benkler, 1998, 2001, 2006; Bollier, 2008). Lessig (1996), for instance, writing in that period about the future of cyberspace, gave the following description:

> As it is just now, cyberspace is such a place of relative freedom. The technologies of control are relatively crude. Not that there is no control. Cyberspace is not anarchy. But that control is exercised through the ordinary tools of human regulation - through social norms, and social stigma; through peer pressure, and reward. How this happens is an amazing question - how people who need never

meet can establish and enforce a rich set of social norms is a question that will push theories of social norm development far. But no one who has lived any part of her life in this space as it is just now can doubt that this is a space filled with community, and with the freedom that the imperfections of community allows. (p. 1407)

In line with the content of Remark 2, most of this literature tends to relate the emergence of such a loosely controlled social space in the non-commercial network with the poor initial development of control-enhancing technologies. In the early days of cyberspace (1970s and 1980s), in fact, control supportive tools such as DRM[12] and DPI systems[13] were not as fully developed as they are today, and were thus easy targets of users' hacking (i.e. low γ) (Zittrain, 2008). At the same time, the low malleability of digital technologies (i.e. high δ) made it relatively costly for platform designers to use code as an effective instrument of regulation. The combination of these two factors created a technical environment that was highly conductive to the emergence of an E_1-type of equilibrium (i.e. it increased the latter's basin of attraction), which was indeed the final state to which the economy converged.

With respect to this interpretation, the above model adds two important points. First of all, the model makes clear that the emergence of a social space characterized by relatively little control was only one of the possible ways in which the public Internet could have evolved. During the 1970s and 1980s there are indeed several examples of proprietary networks that evolved along completely different dynamic paths, becoming in the end highly controlled social environments (e.g. CompuServe, The Source, America Online, Prodigy) (Lessig, 2006). According to Zittrain (2008, 2006), and in line with the results of the model, these networks were relatively inefficient as compared to the public Internet, because they were unable to mobilize a sufficiently high degree of users' participation. As as a result they almost disappeared from the landscape of digital communication. Nonetheless, they do represent clear examples of what alternative systems based on tight forms of control could eventually look like,

[12]Digital rights management (DRM) systems are an example of access control technology that adds code to digital content that disables the simple ability to copy or distribute that content - at least without the technical permission of the DRM system itself (Lessig, 2006). Presently, DRM is in common use by the entertainment industry (e.g. audio and video publisher). Many on-line music stores, such as Apple Inc.'s iTunes Store, as well as many e-book publisher also use DRM, as do cable and satellite service operators to prevent unauthorized use of content or services.

[13]Deep packet inspection (DPI) systems are a form of computer network packet filtering that read and classify Internet traffic as it passes through a network, enabling the identification, analysis, blockage and even alteration of information (MacKinnon, 2012). Initially, DPI were used mainly to secure private internal networks. Recently, Internet service providers (ISPs) have also started to apply this technology on the public network provided to consumers. Common uses of DPI by ISPs are lawful intercetp, policy definition and enforcement, targeted advertising, quality service and copyright enforcement.

and thus provide a direct benchmark against which future developments can be compared.

In addition to this, the model suggests that the emergence of a loosely controlled social space in the public Internet was neither the result of pure chance, nor the unavoidable consequence of the high cost of digital control. Rather, it has deep roots in the specific set of norms and values that formed the early culture of the networked environment. As reported in many analyses on the history of the Internet, in fact, the original population of on-line users consisted for the most part of academics and amateurs who looked at the emerging Internet infrastructure more as an instrument to enhance the human capabilities to communicate and share knowledge, rather than as a tool for running business (Leiner et al., 2001; Berners-Lee, 1999; Abbate, 1999; Wu, 2010). Most of these users exhibited strong intrinsic motives for their on-line actions, and behaved according to a well defined set of ethical norms (Zittrain, 2008; Himanen, 2001; Sterling, 2002). This contributed to generate a cultural environment (i.e. an initial distribution of behavioral types) in which designers employing week forms of control tended to perform far better than those exercising high control, because they were better capable of taking advantage of users' motivation while at the same time saving on the costs of control. Overtime, the dynamic adaptation of digital designs to the cultural features of surrounding environment led to the convergence towards an E_1-type of equilibrium, with low control practices becoming largely predominant. In this sense, both the evidence and the model suggest that the public Internet bore from the very beginning the "cultural seeds" that were necessary for a loosely controlled social space to actually emerge. Quite interestingly, it is exactly the composition of such cultural seeds that got completely overturned as soon as the network was opened to commerce.

4 Opening the network to commerce

The opening of the network to commercial uses occurring in 1995 brought two main changes. First of all, it dramatically increased the size of the populations of both on-line users and designers. Secondly, it brought an upsurge in the number of security incidents associated with attacks to Internet-connected systems (e.g. diffusion of viruses, worms and spams). These two effects are well captured by the data reported in Figures 3 and 4. Figure 3 shows ISC's data[14] on the evolution in the number of Internet hosts during the period 1982-2012.

[14]Internet Systems Consortium (ISC) is a non-profit public benefit corporation dedicated to supporting the infrastructure of the universal connected self-organizing Internet - and the autonomy of its participants - by developing and maintaining core production quality software, protocols, and operations. For more detail on ISC and the data reported in Figure 3 see http://www.isc.org (last time checked: April 30, 2012).

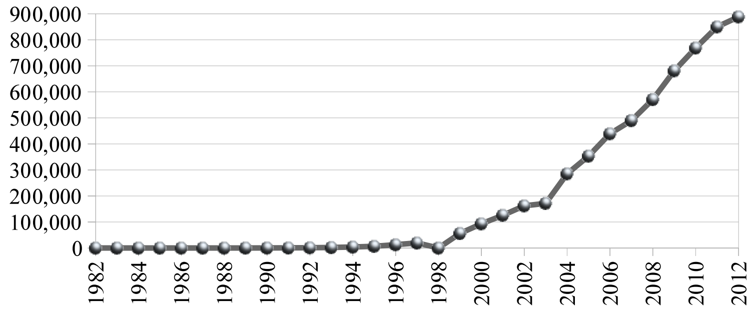

Figure 3: Thousands of Internet hosts, 1982-2012. Source: The Internet System Consortium Domain Survey - Internet host count history, 1981-present.

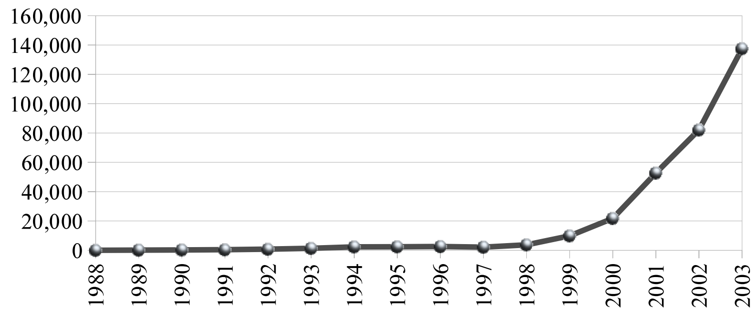

Figure 4: Number of security incidents reported to CERT, 1988-2003. Source: CERT Coordination Center, CERT/CC Statistics 1988-2005.

Figure 4 presents figures on the trend of security incidents reported to the US Department of Defense's CERT Coordination Center for the sub-period 1988-2003.[15] The two graphs show that, starting in 1998, there has indeed been a dramatic increase in the number of both hosts and incidents, with the latter roughly doubling each year through 2003. The two trends, at least for the overlapping time window that I consider, look surprisingly aligned, and there are reasons to believe that a similar tendency extended well beyond 2003. Other informative sources report in fact a constant increase in the rate of Internet vulnerabilities all the way up until the most recent years (see Zittrain, 2008).

The remarkable growth in security incidents that followed the advent of commerce on the Internet has been the subject of several studies. Zittrain (2006, 2008), in particular, links it to the massive increase in the number of un-

[15]The Computer Emergency Response Team (CERT) Coordination Center is a research center located at Carnegie Mellon University's Software Engineering Institute with the aim of studying Internet security vulnerabilities. The same data were originally reported by Zittrain (2006). The data are available only for the period 1988-2003 because in 2004 CERT announced it would no longer keep track of security incidents, since attacks had become so commonplace to be indistinguishable from one another.

skilled and inexpert users, who became easy targets of malware developers and spammers.[16] What is certainly true is that the reduced security of Internet connections is a symptom of a deep cultural change that took place in cyberspace starting in 1995. As the network became so ubiquitous, in fact, the Net-wide set of ethics that worked so well in sustaining the quality of on-line transactions under the loosely controlled environment of the pre-commercial era began to waver. A large number of new users (and designers) who were relatively unused to ethics of cyberspace started to enter in the digital space, causing a significant change in the distribution of behavioral types. Purely extrinsic motives became an important driver behind users' on-line actions (e.g. diffusion of e-commerce), to an extent that ethical values started to be quite often subdued to the possibility of earning monetary rewards. In this sense, the rising business model backing the diffusion of viruses and malware can be seen as a direct consequence of this type of change (Zittrain, 2008).

In order to formally investigate the effect of such change on the equilibrium selection dynamics presented in the previous section, I follows two steps. First of all, I assume the existence of an outside population of users and designers that each period are randomly selected in subsamples to enter the digital economy. Secondly, I study how the distribution of types in these subsamples influences the probability that a transition to a different type of cultural-technological equilibrium occurs. The key assumption that I introduce is that, for any τ, the distribution types in the outside population is independent of the distribution of types in the inside population. This allows me to transform the economy in a stochastic environment, with the distribution of types changing over time for both endogenous and exogenous reasons. Given this framework, I am interested in identifying the conditions under which each of the two cultural-technological equilibria qualifies as the stochastically stable state of the economy.

From the technical point of view I proceed as follows. I call $s_u^\tau = s_I^\tau + s_E^\tau$ (> 0) and $s_d^\tau = s_L^\tau + s_H^\tau$ (> 0) the subsample of users and designers that in each period τ may be selected to enter the economy, with s_i (for $i = \{I, E, L, H\}$) being the number of agents (users and designers) of type i. On this basis, I define:

$$\nu_I^\tau = \frac{s_I^\tau}{s_I^\tau + s_E^\tau} \quad \nu_L^\tau = \frac{s_L^\tau}{s_H^\tau + s_L^\tau} \tag{16}$$

as the fractions of I-users and L-designers existing in this subsample. As previously stated, I assume the value of ν_I^τ and ν_L^τ to be independent of ω_I^τ and ω_L^τ at any τ. In particular, I assume the former to be random draws from the probability distributions $f_I(\nu)$ and $f_L(\nu)$, with $f(x)$ continuous over the interval $x \in [0, 1]$. I call $\overline{\nu}_I = \int_{-\infty}^{+\infty} \nu f_I(\nu) \mathrm{d}\nu$ and $\overline{\nu}_L = \int_{-\infty}^{+\infty} \nu f_L(\nu) \mathrm{d}\nu$ the expected

[16] Software designed to infiltrate and damage a computer system (Zittrain, 2006).

fractions of I-users and L-designers in the selected subsample. The latter can be indeed seen as indexes of how homogeneous the distribution of types is in the outside population.

The timing of entrance is modeled as follows. At the beginning of any period τ, users and designers update their type following the process described in the previous section. This process takes as a reference the distribution of preferences and designs that exist in the economy at the beginning of that period. Once such updating is completed, Nature makes two moves. First, she determines with probability $\epsilon d\tau$ (> 0) whether a new set of s_u users and s_d designers enter the economy. Second, she selects the value of ν_I^τ and ν_L^τ to be associated to that set. For the sake of simplicity I assume that both populations of users and designers grow at the same constant rate so that, for any τ, we have $s_u^\tau = \rho\, n_u^\tau$ and $s_d^\tau = \rho\, n_d^\tau$ with $\rho > 0$.

The effective possibility of external entrance transforms the dynamical system into an ergodic process, with transitions between the basins of attraction of the two equilibria E_1 and E_0 that now becomes possible. Whenever a new entrance occurs, in fact, the distribution of types at any given τ reflects both the endogenous updating undertaken by the inside population and the exogenous variation due to the new entrants. When the effect of the latter is sufficiently strong, the population can be drifted away from the status-quo convention, and eventually converge towards a new equilibrium. In order to see why, let us suppose that the population is in equilibrium E_1 and entrance occurs. Once the individual updating is completed, the fraction of I-users and L-designers at the beginning of next period can be written as follows:

$$\omega_I^{\tau+1} = \frac{1 + \rho\nu_I^\tau}{1 + \rho} \qquad \omega_L^{\tau+1} = \frac{1 + \rho\nu_L^\tau}{1 + \rho} \tag{17}$$

where I used the fact that $\omega_I^\tau = 1$, $\omega_L^\tau = 1$, $s_u^\tau = \rho\ n_u^\tau$ and $s_d^\tau = \rho\ n_d^\tau$. A transition from E_1 to E_0's basin of attraction will occur whenever $1 - \omega_I^{\tau+1} \leq 1 - \omega_I^*$ and/or $1 - \omega_L^{\tau+1} \leq 1 - \omega_L^*$, which is the case if

$$\nu_I^\tau \leq \frac{\omega_I^*(1+\rho) - 1}{\rho} \quad \textit{and} \quad \nu_L^\tau \leq \frac{\omega_L^*(1+\rho) - 1}{\rho} \tag{18}$$

Depending on the value of ρ and the shape of $f(\nu)$, therefore, the "tipping" of the population from one basin of attraction to the other is more or less likely to occur.

Figure 5 offers a graphical representation of the way in which motivation (upper panel) and control (lower panel) may co-evolve in this stochastic environment (for the derivation of the underlying dynamical system see Appendix B). The black and gray lines represent two distinct runs of 500 iterations, with

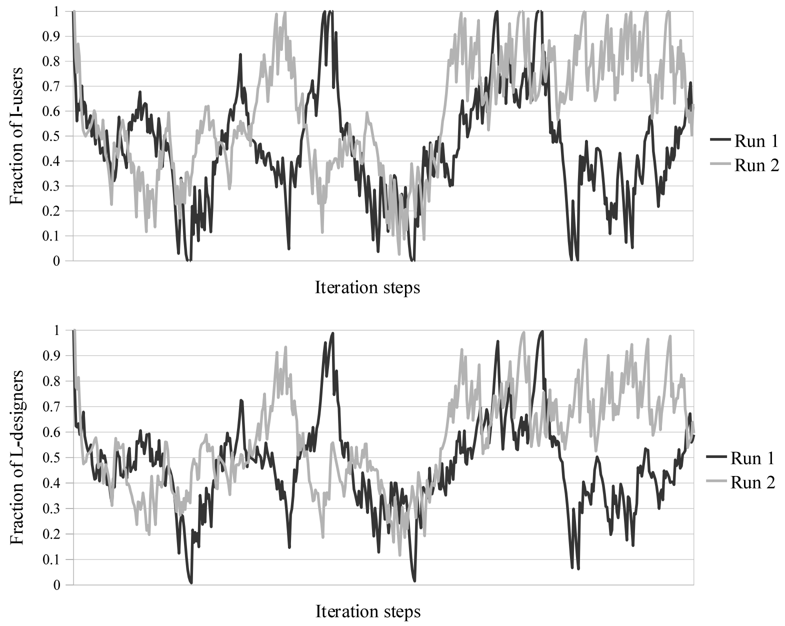

Figure 5: Evolution of motivation and control in a stochastic environment. Note: $\phi = 2.4$, $\lambda = 1$, $q = 1$, $\gamma = 1$, $\eta = 0.33$, $k = 2.41$, $\mu = 2.91$, $\delta = 2.21$, $\alpha\beta = 0.4$, ν_I^τ and ν_I^τ are random draws from the uniform distribution $[0.3, 0.7]$.

starting point at $\{1, 1\}$. The system is calibrated using the parameters reported the Figure's caption. In particular, I assume $f_I(\nu)$ and $f_L(\nu)$ to be a uniform distribution over the interval $[0.3, 0.7]$, with $\overline{\nu}_I = \overline{\nu}_L = 0.5$. As it is easy to see the evolution of individual types in the two populations follows a closely related path. In Run 1 such path oscillates between the two basins of attraction for all 500 iterations. In Run 2 the dynamics oscillates too for nearly half of the iterations, and then it tends stabilize in the orbit of equilibrium E_1. In the two cases, starting from the same initial conditions, the population follows two completely different dynamic paths with transitions between the two stable equilibria being relatively frequent.

Among the several factors that may explain both the speed and frequency of transitions, one that appear to be of major relevance for the present discussion concerns the distribution of types in the population of new entrants. The more such distribution is biased towards the predominance of one particular type, in fact, the more the exogenous variation will tend to keep the economy close to one specific equilibrium and make transition in the opposite direction unlikely to occur. Quite interestingly, this is true also when such bias concerns only one of the two populations of new entrants, being either users and designers.

On this respect, Figure 6 shows the evolution of control when the distribution of behavioral types in the population of new users is biased in favor of

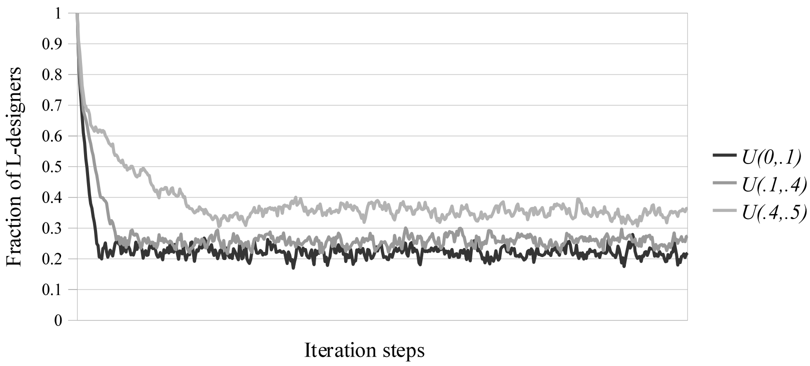

Figure 6: Emergence of control for different distribution of behavioral types in the outside population. Note: each curve is the average of 40 runs under the indicated uniform distribution; all other parameters are as in Figure 5.

E-type (motivation obviously follows a closely related path). Each curve represents the average of 40 runs under the indicated uniform distribution. All the other parameters are kept the same as in Figure 5. As it is easy to observe, small variations in the distribution of behavioral types significantly change the evolution of control. In all cases, the fraction of H-designers tends to increase over time and become largely predominant. The more the distribution is biased in favor of E-type (from the lightest to darkest curve), the faster the convergence towards a high-control-type of equilibrium and the greater the fraction of H-designers in the stable path.

In addition to the distribution of individual types, another factor that plays a crucial role in influencing the shape of the overall dynamics concerns the two critical values ω_I^* and ω_L^*. As reported in Eq. (18), in fact, the latter contribute to the definition of the threshold values against which a transition between the two basins of attraction is made possible, and therefore affect the amount of exogenous variation that is actually necessary for such a switch to occurs. Intuition suggests that the cultural-technological equilibrium that requires more exogenous variation to dislodge, and less exogenous variation to access will tend to persist longer than the other. At the same time, if dislodged, it will tend to reemerge readily. This, at least for a sufficiently homogeneous distribution of types in the outside population, is the cultural-technological equilibrium that is most likely to be observed in the long-run.

In order to formalize the above intuition consider the following definitions - both adaptations from Young (1998) and Bowles (2006):

Definition 3. *Let r_{jk}, the reduced resistance on the path from E_j to E_k, be* *the minimal fraction of agents (users and designers) that, should the population's type frequencies after entrance be greater or equal r_{jk}, would induce the best-responding partners to switch their types. Then,* $r_{01} = \min\{\omega_I^*, \omega_L^*\}$ *and* $r_{10} = \min\{1 - \omega_I^*, 1 - \omega_L^*\}$.

Definition 4. *The stochastically stable equilibrium (SSE) is the one that occurs with non-negligible probability when the rate of exogenous variation is arbitrarily small. In a* 2×2 *coordination game with two asymptotically stable equilibria* E_j *and* E_k, E_j *is SSE if and only if* $r_{kj} < r_{jk}$.

Definitions 3 and 4 can be used in order to find the conditions under which each of the two asymptotically stable cultural-technological equilibria identified in Proposition 3 qualifies as SSE. This turns out to be of particular interest especially if related to the efficiency properties of the two equilibria (see Proposition 2). In particular, I obtain the following result :

Proposition 4. *Suppose* $\epsilon > 0$. *Then, in the dynamic system with exogenous variation there exist a* $k^* = [\psi(2q\lambda+\delta)+2\mu\delta]/2\gamma(2\mu+\eta\psi)$ *where* $\psi = \lambda(\lambda+2\phi)$ *such that if* $k < k^*$ E_0 *is SSE. This is true even if* E_0 *is Pareto inefficient.*

The intuition behind Proposition 4 is straightforward. When the cost associated with the realization of a noisy tasks k is sufficiently high, L-designers suffers a big loss whenever they are matched with an E-users (see payoffs in Table 1). This implies that when there is uncertainty concerning the distribution of types among users - because for instance there is a positive rate of exogenous variation, a H-type design will tend to have a selection advantage over an L-type, because it ensures a greater expected payoff. This amounts to say that offering a H-type design is risk-dominant in the standard sense that if one believes that users are either I-type or E-type with equal probability, then the best response is to offer a H-type design.[17] Over time, H-designers will tend therefore to increase in number causing a contemporaneous reduction in the number of I-agents. The greater the average degree of control in the economy (i.e. the larger the fraction of H-designers), in fact, the stronger the crowding out effect on motivation, and thus the larger the number of agents who become purely extrinsically motivated. The combination of these effects make equilibrium E_0 stochastically stable.

The content of Proposition 4 has interesting implications for what concerns the evolution of digital control. In spite of the efficiency advantage of equilib-

[17] On the relationship between risk-dominance and stochastic stability see Foster and Peyton Young (1990).

rium E_1, in fact, I find that there exist a whole range of values in the parameter space which makes equilibrium E_0 persistent over time. Whether the economy is actually in (or is likely to converge to) this state is impossible to say theoretically, and becomes mainly an empirical question. What the model shows is just the possibility that such Pareto inefficient state may become the long-run cultural-technological equilibrium of the digital space. This in turn calls for a serious analysis of the policy regime that is currently governing cyberspace, with particular attention on some recent proposals for Internet regulation.

5 Policy implications

The results of the previous sections depict the possibility of a coordination failure in the evolution of digital control. When cost k is sufficiently high, tight forms of control and extrinsic motivation tend to become predominant in the economy, leading to the persistence of equilibrium E_0. In some cases this outcome is suboptimal from the social point of view, because both users and designers would be better-off if they could only coordinate their actions in favor of equilibrium E_1. When this happens, the economy is trapped in a low efficiency equilibrium and some forms of government intervention can be justified.

Although it is probably too early to say whether such a coordination failure will effectively emerge, it is still possible to analyze the effect of different forms of government intervention in increasing and/or reducing the likelihood of its occurrence. On this respect, several policy proposals that have been recently discussed at both the national and international level seem to be of relevance. Three, in particular, have attracted the attention of most international commentators, and include ACTA, the SOPA and PIPA bills in the U.S., and the Google-Verizon's proposal on network neutrality. Although none of these proposals is directly concerned with the implementation of control *per se*, they all impact on the latter's effectiveness and appropriateness. As a result, they all directly affect the probability that a coordination failure of the type described above may effectively occur.

ACTA is a proposed multinational treaty that aim at establishing international standards for intellectual property rights enforcement. According to the original proposing parties, the main objective of ACTA is to help fight the proliferation of counterfeit and pirated goods in international trade, which have by now become one major source of profit for illegal and criminal activities, especially in developing countries (McManis, 2008). If finally approved, ACTA would apply 'new, stricter legal and enforcement standards to the trade in informational goods', and introduce 'sweeping provisions to criminalize information use practices currently allowed under U.S., European and international law' (Shaw, 2008, p.1). With specific reference to trade in digital goods, ACTA aims at 'reinforcing so-called "Digital Rights Management" (DRM) technologies that currently prevent the personal, legal reproduction of optical discs like DVDs and trample on "fair use" rights'. In addition, it proposes to 'undermine legal safeguards that protect Internet Service Providers (ISPs) from the liability of the actions of their subscribers' (Shaw, 2008, p.3). In the language of the above model, the implementation of these provisions would at the same time improve the effectiveness of control - via the strengthening of DRM systems, i.e. increase in γ, and rise the state-enforced costs for copyright infringement - via the increased liability of ISPs, i.e. increase in k. The combination of these two effects would make digital control increasingly convenient as a design option, thus favoring the convergence towards equilibrium E_0. If that happens, the undeniable benefit that is associated with a reduction of illegal and criminal activities, would be then counterbalanced by the increased risk of altering the cultural-technological features of the digital space in a socially inconvenient way. This would in turn question the effective appropriateness and applicability of the treaty itself.

A very similar interpretation holds also for another set of legal provisions that explicitly aim at fighting the problem of on-line piracy, such as the SOPA and PIPA bills under scrutiny in the U.S. Congress. The two bills, introduced respectively in the House and the Senate, are the most recent iteration of the long list of acts aimed at strengthening the rights of the U.S. copyright industries, such as the Digital Millennium Copyright Act of 1998, the Prioritizing Resources and Organization for Intellectual Priority Act of 2008, the Higher Education Opportunity Act of 2008, and the Combating On-line Infringement and Counterfeits Act of 2010. In their current version, the bills' provisions aims at further extending the involvement of criminal enforcement authorities in what was traditionally an area of private commercial law, and at using the state leverage to harness private platform providers to enforce the interests of copyright holders (Benkler, 2012b). Similarly to ACTA, the SOPA and PIPA bills intend to curb criminal and illegal on-line practices, while at the same ensuring the defense of individual rights. In doing so, however, they create an environment in which control becomes at the same time easy to implement (low γ) and costly to avoid (high k), thus making equilibrium E_0 likely to emerge. Whether the cost of this relatively inefficient outcome is compensated by the benefit associated with a reduced degree of on-line piracy is difficult to say, and requires an in-depth empirical investigation. What is certainly true is that simple possibility of such trade-off suggests the need of a partial rethinking of the bills' content, with particular attention on the role played by the Internet's traditional openness.

Finally, a third type of intervention that can directly affect the evolution of digital control is the Google-Verizon's proposal on network neutrality. The proposal, presented to the U.S. Federal Communication Commission on August 2010, introduces the possibility to exempt wireless communication and other on-line services from applying the principle of net neutrality, that is the set of embedded rules which impose that all like Internet content must be treated alike and move at the same speed over the network (Wu, 2003a). The proposal finds its rationale in the conviction that by allowing ISPs to (at least partially) discriminate on some Internet applications, new economic resources could be generated, which could be in turn invested in the creation of new, more efficient broadband and information technology services[18]. Some commentators, however, urges that a similar provision would at the same time increase the discretionary power of ISPs, making it possible to tailor specific types of code-based restrictions on Internet applications.[19] With reference to the model, this would imply a substantial reduction in the design cost of digital control (i.e. δ), and thus an increase in the persistence of equilibrium E_0. Whenever the latter is Pareto inefficient, this would in turn generate a trade-off between the provision of incentives to invest in innovation and the distortion of the cultural-technological features of the digital space. Once again, a sound balance between these types of costs and benefits is effectively difficult to strike.

Overall, the analysis of three of the most recent policy proposals for Internet regulation reveals a relatively complex scenario. If on one hand the proposed interventions pursue fairly legitimate policy objectives, on the other they all introduce provisions that tend to increase the chances that the economy gets stuck in a low efficiency equilibrium. The reason is essentially related to the fact that, while being concerned with the enforcement of particular rights and the creation of specific incentives, these laws tend to neglects the economy-wide effects that an increased viability of digital control may have on the cultural and technological features of the networked environment. Whenever these effects are worse than the benefits the laws are aimed at generating, the policy prescriptions should be revised, and the preservation of the cultural-technological features of cyberspace should become an integral part of policy design.

[18]See Google-Verizon Proposal for a legislative framework for network neutrality, available at: `http://static.googleusercontent.com/external_content/untrusted_dlcp/www.google.com/it//googleblogs/pdfs/verizon_google_legislative_framework_proposal_081010.pdf` (last time checked: May 3, 2012).

[19]See Cain Miller and Helft, supra note 5.

6 Conclusion

On August 2011, while speaking with The Associated Press on the sidelines of the 7th Wikipedia's annual conference, Jimmy Wales (the website's founder) said the on-line encyclopedia was struggling to find contributors. After many years of constant growth, the non-profit organization reported that contributors leaving the website had outnumbered new users, leaving the community in short supply. Although Wales linked this poor result to the website's complex editing procedures,[20] it can be interpreted as the first sign of a deeper change. In a highly competitive environment that is increasingly populated by control-intensive platforms (e.g. social networks), in fact, open and commons-based websites like Wikipedia finds it increasingly difficult to attract deeply motivated users. Whether this implies that such kind of platforms are effectively doomed to disappear it is difficult to say; but this evidence certainly suggests that something is changing in the way in which on-line participation is being experienced.

Starting from this evidence, this paper has presented a behavioral economic model that micro-founds the cultural-technological evolution of the digital space. The paper focused on the interaction between individual motivation and digital control, with the aim of modeling the latter's historical evolution. The crucial assumption of the model was that control is not neutral with respect to the nature of individual motivation, and some crowding out effects on intrinsic motives exists. On this basis, the paper has derived three main results: a) in the long-run there exist two stable cultural-technological equilibria in the digital economy: one with intrinsically motivated users and low control; and the other with purely extrinsically motivated users and high control; b) under a closed economy - i.e. before the opening of the network to commerce, the initial emergence of a low-control-intrinsic-motivation equilibrium can be explained by the specific set of norms and values that formed the early culture of the networked environment; and c) the opening of the network to commerce can indeed cause a transition to a high-control-extrinsic-motivation equilibrium, even if the latter is Pareto inferior. Although it is too early to say whether such a transition is actually taking place, these results call for a great deal of attention in evaluating policy proposals on Internet regulation. This paper, in particular, has focused on three of them, such as ACTA, the SOPA and PIPA bills, and the Google-Verizon's proposal on network neutrality. All of them are deemed to have controversial effects on the cultural-technological nature of cyberspace.

[20] See Zack Whittacker, *Wikipedia losing contributors: Fatal flaw, the community editors?*, ZDNet, Augist 4, 2011, available at: `http://www.zdnet.com/blog/btl/wikipedia-losing-contributors-fatal-flaw-the-community-editors/54144` (last time checked: May 3, 2012)

A Appendix

A.1 Proof of Lemma 1.

The derivative of Eqs. (4) with respect to a gives us the following best-response function for I- and E-users when paired with a generic designer j: $a_{I,j} = \phi + \lambda(1-t)$ and $a_{E,j} = \phi$. By substituting away for t we obtain the best-response level of a reported in the lemma.

A.2 Proof of Proposition 1

$\{I, L\}$ is proven to be Nash equilibrium as long as: (a) $(phi + \lambda)^2/2 > \phi^2/2$, and (b) $q(\phi + \lambda) - \gamma\eta k > q\phi - \delta/$. Condition (a) is self-explained. Condition (b) reduces to $\delta > 2(\gamma\eta k - q\lambda) = \underline{\delta}$. Similarly, $\{E, H\}$ is a Nash equilibrium as long as: (c) $\phi^2/2 > \phi^2/2 - \mu$ and (d) $q\phi - \gamma k < q\phi - \delta/2$. Condition (c) is self-explained. Condition (d) reduces to $\delta < 2\gamma k = \overline{\delta}$. For $0 < \eta < 1$, $\underline{\delta} < \overline{\delta}$ is always true. It follows that: *(i)* when $\delta > \overline{\delta}$ condition (b) is satisfied but not condition (d), hence $\{I, L\}$ is the only Nash equilibrium; *(ii)* when $\delta < \underline{\delta}$ condition (d) is satisfied but not condition (b), hence $\{E, H\}$ is the only Nash equilibrium; and (iii) when $underline\delta < \delta < \overline{\delta}$ conditions (b) and (d) are simultaneously satisfied, hence both $\{I, L\}$ and $\{E, H\}$ are Nash equilibria. Corollary 1.1 follows from the fact that two necessary conditions for $\{E, L\}$ and $\{I, H\}$ to be Nash equilibria are that E is a best-response to L and I is a best response to H, but this is impossible because it would violate conditions (a) and (c) above. Corollary 1.2 follows directly from points (i), (ii) and (iii) above.

A.3 Proof of Proposition 2.

For any $\lambda > 0$, a necessary and sufficient condition for $\{E, H\}$ to be Pareto efficient is that $q(\phi + \lambda) - \gamma\eta k < q\phi - \delta/$, which reduces to $\delta < \overline{\delta}$. Otherwise, $\{I, L\}$ Pareto dominates $\{E, H\}$. This, together with the results of Proposition 1, implies that: (i) if $\delta < \overline{\delta}$, then $\{E, H\}$ is Pareto efficient and it is also the only Nash equilibrium of the game; (ii) if $\delta > \overline{\delta}$, then $\{I, L\}$ is Pareto dominant and it is also a Nash equilibrium. Points (i) and (ii), together with the fact that for $\underline{\delta} < \delta < \overline{\delta}$ both $\{I, L\}$ and $\{E, H\}$ are Nash equilibria, prove the proposition.

A.4 Proof of Proposition 3.

The five cultural-technological equilibria are derived by simply solving the system (13)-(14) for $\dot{\omega_\tau} = 0$ and $\dot{\rho_\tau} = 0$. The proof in this case is omitted. The asymptotic properties of each equilibrium are derived by analyzing the Jacobean

Matrix $J(\omega_I, \omega_L)$ associated to system (13)-(14), which takes the following form:

$$J = \begin{pmatrix} (1-2\omega_I)\left[\omega_L\left(\frac{\lambda^2}{2}+\phi\lambda+\mu\right)-\mu\right] & (\omega_I-\omega_I^2)\left(\frac{\lambda^2}{2}+\phi\lambda+\mu\right) \\ (\omega_L-\omega_L^2)\left[q\lambda+\gamma k(1-\eta)\right] & (1-2\omega_L)\left\{\omega_I\left[q\lambda+\gamma k(1-\eta)\right]+\frac{\delta}{2}-\gamma k\right\} \end{pmatrix}$$

At $\{0,0\}$, we have

$$J = \begin{pmatrix} -\mu & 0 \\ 0 & \frac{\delta}{2}-\gamma k \end{pmatrix}$$

from which it follows that

$$Tr(J) = -\mu + \frac{\delta}{2} - \gamma k \qquad and \qquad Det(J) = -\mu\left(\frac{\delta}{2}-\gamma k\right) \tag{19}$$

Since $\mathrm{Tr}(J) < 0$ and $\mathrm{Det}(J) > 0$ for any $\delta < 2\gamma k$, $\{0,0\}$ is asymptotically stable.

At $\{1,0\}$, we have

$$J = \begin{pmatrix} \mu & 0 \\ 0 & q\lambda - \gamma\eta k + \frac{\delta}{2} \end{pmatrix}$$

from which it follows that

$$Tr(J) = \mu + q\lambda - \gamma\eta k + \frac{\delta}{2} \qquad and \qquad Det(J) = \mu\left(q\lambda - \gamma\eta k + \frac{\delta}{2}\right) \tag{20}$$

Since $\mathrm{Tr}(J) > 0$ and $\mathrm{Det}(J) > 0$ for any $\delta > 2(\gamma\eta k - q\lambda)$, $\{1,0\}$ is unstable.

At $\{0,1\}$, we have

$$J = \begin{pmatrix} \frac{\lambda^2}{2}+\phi\lambda & 0 \\ 0 & -\frac{\delta}{2}+\gamma k \end{pmatrix}$$

from which it follows that

$$Tr(J) = \frac{\lambda^2}{2}+\phi\lambda-\frac{\delta}{2}+\gamma k \qquad and \qquad Det(J) = \left(\frac{\lambda^2}{2}+\phi\lambda\right)\left(\gamma k - \frac{\delta}{2}\right) \tag{21}$$

Since $\mathrm{Tr}(J) > 0$ and $\mathrm{Det}(J) > 0$ for any $\delta < 2\gamma k$, $\{0,1\}$ is unstable.

At $\{1,1\}$, we have

$$J = \begin{pmatrix} -\frac{\lambda^2}{2} - \phi\lambda & 0 \\ 0 & -q\lambda + \gamma\eta k - \frac{\delta}{2} \end{pmatrix}$$

from which it follows that

$$\begin{aligned} Tr(J) &= -\frac{\lambda^2}{2} - \phi\lambda - q\lambda + \gamma\eta k - \frac{\delta}{2} \qquad and \\ Det(J) &= -\left(\frac{\lambda^2}{2} + \phi\lambda\right)\left(\gamma\eta k - q\lambda - \frac{\delta}{2}\right) \end{aligned} \tag{22}$$

Since Tr(J) < 0 and Det(J) > 0 for any $\delta > 2(\gamma\eta k - q\lambda)$, $\{1, 1\}$ is asymptotically stable.

At $\{\omega_I^*, \omega_L^*\}$, we have

$$J = \begin{pmatrix} 0 & \frac{(2\gamma k - \delta)\,[2q\lambda - 2\gamma\eta k + \delta]}{4\,[q\lambda + \gamma k(1-\eta)]^2}\left(\frac{\lambda^2}{2} + \phi\lambda + \mu\right) \\ \frac{2\mu\lambda(\lambda + 2\phi)}{[\lambda(\lambda + 2\phi) + 2\mu]^2}\,[q\lambda + \gamma k(1-\eta)] & 0 \end{pmatrix}$$

from which it follows that

$$\begin{aligned} Det(J) = -\frac{(2\gamma k - \delta)\,[2q\lambda - 2\gamma\eta k + \delta]}{4\,[q\lambda + \gamma k(1-\eta)]^2}\left(\frac{\lambda^2}{2} + \phi\lambda + \mu\right) \cdot \\ \cdot \frac{2\mu\lambda(\lambda + 2\phi)}{[\lambda(\lambda + 2\phi) + 2\mu]^2}\,[q\lambda + \gamma k(1-\eta)] \end{aligned} \tag{23}$$

Since Det(J) < 0 for any $\delta > 2(\gamma\eta k - q\lambda)$, $\{\omega_I^*, \omega_L^*\}$ is a saddle.

A.5 Proof of Proposition 4.

From Definition 3 and the value of ω_I^* and ω_L^* reported in Proposition 3 it follows that:

- $\mu \geq \psi(2\gamma k - \delta)/2\,[\delta + 2(q\lambda - \gamma k\eta]$ $\quad\Longleftrightarrow$

$$r_{01} = \omega_I^* = \frac{2\gamma k - \delta}{2\,[q\lambda + \gamma k(1-\eta)]} \quad and \quad r_{10} = 1 - \omega_L^* = \frac{\psi}{\psi + 2\mu} \tag{24}$$

- $\mu < \psi(2\gamma k - \delta)/2\,[\delta + 2(q\lambda - \gamma k\eta]$ $\quad\Longleftrightarrow$

$$r_{01} = \omega_L^* = \frac{2\mu}{\psi + 2\mu} \quad and \quad r_{10} = 1 - \omega_I^* = \frac{\delta + 2(q\lambda - \gamma k\eta)}{2\,[q\lambda + \gamma k(1-\eta)]} \tag{25}$$

where $\psi = \lambda(\lambda + 2\phi)$. According to Definition 5, E_0 is SSE if and only if $r_{10} < r_{01}$. Simple algebra shows that, given Eqs. (24) and (25), the latter condition holds if and only if $k < [\psi(2q\lambda+\delta)+2\mu\delta]/2\gamma(2\mu+\eta\psi) = k^*$. The second part of the proposition follows directly from Proposition 2.

B Appendix

B.1 Payoffs in Table 1

Let us indicate with $U_{i,j}$ the utility of an i-type user when matched with a j-type designer, and with $\pi_{j,i}$ the return to an j-type designer when matched with a i-type user. Moreover let us write $a_{i,j}$ as the best-response level of a for an i-type user when matched with a j-type designer. Given Eqs. (1), (2) and the functional forms defined in Section 3.1, we have:

$$U_{I,j} = [\phi + \lambda(1-t)]a_{I,j} - \frac{a_{I,j}^2}{2} - \mu t \quad , \quad U_{E,j} = \phi a_{E,j} - \frac{a_{E,j}^2}{2} \tag{26}$$

$$\pi_{L,i} = q a_{i,L} - \gamma\eta(\lambda)k \quad , \quad \pi_{H,i} = q a_{i,H} - \frac{\delta}{2} \tag{27}$$

where $\eta(\lambda)$ takes the following form:

$$\eta(\lambda) = \begin{cases} 1, & \text{if } i = E \\ \\ \eta, & \text{if } i = I \end{cases} \tag{28}$$

By replacing into Eqs. (26) and (27) the value for $a_{i,j}$ reported in Lemma 1, and substituting away for t (i.e. replacing $t = 0$ and $t = 1$ for a match with an L- and a H-type designer respectively), we obtain the following results:

$$U_{I,L} = \frac{(\phi+\lambda)^2}{2} \quad , \quad U_{I,H} = \frac{\phi^2}{2} - \mu \quad , \quad U_{E,L} = U_{E,H} = \frac{\phi^2}{2} \tag{29}$$

$$\pi_{L,I} = q(\phi+\lambda) - \gamma\eta k \quad , \quad \pi_{L,E} = q\phi - \gamma k \quad , \quad \pi_{H,I} = \pi_{H,E} = q\phi - \frac{\delta}{2} \tag{30}$$

B.2 Stochastic dynamical system

In the stochastic environment described in Section 4 the expected fraction of I-users in period $\tau + d\tau$ is given by

$$\begin{aligned} \omega_I^{\tau+d\tau} = [\omega_I^\tau - \omega_I^\tau(1-\omega_I^\tau)\alpha d\tau\sigma_E\beta(V_E^\tau - V_I^\tau) + \\ +(1-\omega_I^\tau)\omega_I^\tau\alpha d\tau\sigma_I\beta(V_I^\tau - V_E^\tau)]\chi_u^\tau + \nu_I^\tau(1-\chi_u^\tau) \end{aligned} \tag{31}$$

where

$$\chi_u^\tau = \frac{n_u^\tau}{n_u^\tau + \varepsilon d\tau s_u^\tau} \tag{32}$$

is a normalizing factor that varies according to the number of new users that enter into the economy. The part of Eq. (31) inside the square brackets is the same as Eq. (11) and represents the updating process undertaken by the users that are already part of the economy at the beginning of period τ. Once such updating process is completed, s_u^τ new users enter the economy with probability $\varepsilon d\tau$. The fraction of I-users at the beginning of next period is thus given by the updated fraction of I-users normalized by the new size of the users' population (i.e. multiplication by χ_u^τ), plus the fraction of I-users that are included in the set of new entrants (i.e. $\nu_I^\tau(1-\chi_u^\tau)$). Similarly, the expected fractions of L-designers in period $\tau + d\tau$ is given by:

$$\begin{aligned}\omega_L^{\tau+d\tau} = [&\omega_L^\tau - \omega_L^\tau(1-\omega_L^\tau)\alpha d\tau\sigma_H\beta(V_H^\tau - V_L^\tau)+ \\ &+(1-\omega_L^\tau)\omega_L^\tau\alpha d\tau\sigma_L\beta(V_L^\tau - V_H^\tau)]\chi_d^\tau + \nu_L^\tau(1-\chi_d^\tau)\end{aligned} \tag{33}$$

where

$$\chi_d^\tau = \frac{n_d^\tau}{n_d^\tau + \varepsilon d\tau s_d^\tau} \tag{34}$$

Subtracting ω_I^τ and ω_L^τ from both sides of Eqs. (31) and (33) respectively, dividing both equations by $d\tau$, and taking the limit as $d\tau \to 0$, we get:

$$\dot{\omega}_I^\tau = \omega_I^\tau(1-\omega_I^\tau)(V_I^\tau(\omega_L^\tau) - V_E^\tau(\omega_L^\tau)) + \varepsilon\rho_u(\nu_I^\tau - \omega_I^\tau) \tag{35}$$

$$\dot{\omega}_L^\tau = \omega_L^\tau(1-\omega_L^\tau)(V_L^\tau(\omega_I^\tau) - V_H^\tau(\omega_I^\tau)) + \varepsilon\rho_d(\nu_L^\tau - \omega_L^\tau) \tag{36}$$

where $\rho_u = s_u/n_u$, $\rho_d = s_d/n_d$, and I assumed $\alpha\beta = 1$. Eqs. (35) and (36) represent a system of differential equations which describes how the distribution of types $\{\omega_I^\tau, \omega_L^\tau\}$ evolve over time. The main difference with the system composed of Eqs. (13) and (14) is that this time there are also two stochastic components represented by variables ν_I^τ and ν_L^τ. The latter are the sources of exogenous variation that make a transition between the basins of attraction of the two stable equilibria E_0 and E_1 possible.

References

Abbate, J., 1999. Inventing the Internet. MIT Press, Cambridge.

Belloc, M., Bowles, S., 2011. International trade, factor mobility and the persistence of cultural-institutional diversity, unpublished manuscript.

Benabou, R., Tirole, J., 2003. Intrinsic and extrinsic motivation. Review of Economic Studies 70, 489–520.

Benkler, Y., 1998. Overcoming agoraphobia: Building the commons of the digitally networked environment. Harvard Journal of Law and Technology 11 (2), 287–400.

Benkler, Y., 2001. Siren songs and amish children: Autonomy, information, and law. New York University Law Review 76, 23–113.

Benkler, Y., 2002a. Coase's penguin, or linux and the nature of the firm. The Yale Law Journal 112 (3), 369–446.

Benkler, Y., 2002b. Intellectual property and the organization of information production. International Review of Law and Economics 22, 81–107.

Benkler, Y., 2006. The Wealth of Networks: How Social Production Transforms Markets and Freedom.

Benkler, Y., 2012a. A free irresponsible press: Wikileaks and the battle over the soul of the networked fourth estate. Harvard Civili Rights-Civil Liberties Law Review forthcoming.

Benkler, Y., 2012b. Wikileaks and the protect-ip act: A new public-private threat to the internet commons. Doedalus, the Journal of the American Academy of Arts and Science 140 (4), 154–164.

Benkler, Y., Nissenbaum, H., 2006. Commons-based peer production and virtue. The Journal of Political Philosophy 14 (4), 394–419.

Berners-Lee, T., 1999. Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web. HarperCollins Publisher Inc., New York.

Bisin, A., Verdier, T., 2001. The economics of cultural transmission and the dynamics of preferences. The Journal of Economic Theory 97, 298–319.

Bollier, D., 2008. Viral Spiral: How the Commoners Built a Digital Republic of Their Own. The New Press, New York.

Bowles, S., 1985. The production process in a competitive economy: Walrasian, neo-hobbesian, and marxian. The American Economic Review 75 (1), 16–36.

Bowles, S., 2006. Microeconomics: Behavior, Institutions, and Evolutions. Princeton University Press, Princeton.

Bowles, S., Choi, J.-K., Hopfensitz, A., 2003. The co-evolution of individual behaviors and social institutions. Journal of Theoretical Biology 223, 135–147.

Bowles, S., Hwang, S.-H., 2008. Social preferences and public economics: Mechanism design when social preferences depend on incentives. Journal of Public Economics 92 (8-9), 1811–20.

Bowles, S., Polania-Reyes, S., 2012. Economic incentives and social preferences. Journal of Economic Literature forthcoming.

Charness, G., Cobo-Reyes, R., Jimenez, N., Lacomba, J. A., Lagos, F., 2011. The hidden advantage of delegation: Pareto-improvements in a gift-exchange game. The American Economic Review forthcoming.

Deci, E. L., Ryan, R. M., 1985. Intrinsic Motivation and Self-determination in Human Behavior. Plenum Press, New York.

Deibert, R., Palfrey, J., Rohozinski, R., Zittrain, J. (Eds.), 2010. Access Controlled: The Shaping of Power, Rights, and Rule in Cyberspace. MIT Press, Cambridge.

Elkin-Loren, N., Salzberger, E. M., 2000. Law and economics in cyberspace. International Review of Law and Economics 19, 553–581.

Falk, A., Kosfeld, M., 2006. The hidden costs of control. American Economic Review 96 (5), 1611–30.

Fehr, E., Holger, H., Wilkening, T., 2010. The lure of authority: Motivation and incentice effects of power, unpublished manuscript.

Foster, D. P., Peyton Young, H., 1990. Stochastic evolutionary game dynamics. Theoretical Population Biology 38 (2), 219–232.

Frey, B. S., 1997. Not Hust for the Money: An Economic Theory of Personal Motivation. Edward Elgar Publishing, Inc., Chelthenham, UK.

Frey, B. S., Jegen, R., 2001. Motivation crowding theory: A survey of empirical evidence. The Journal of Economic Surveys 15 (5), 589–611.

Gagne, M., Deci, E. L., 2005. Self-determination theory and work motivation. Journal of Organizational Behavior 26, 331–362.

Goldsmith, J., Wu, T., 2006. Who Controls the Internet? Illusions of a Borderless Wolrd. Oxford University Press, New York.

Himanen, P., 2001. The Hacker Ethic: A Radical Approach to th Philosophy of Business. Random House Inc., New York.

Irlenbusch, B., Ruchala, G. K., 2008. Relative rewards within team-based compensation. Labour Economics 15, 141–167.

Johnson, D. R., Post, D., 1996. Law and borders: The rise of law in cyberspace. Stanford Law Review 48 (5), 1367–1402.

Landini, F., 2012. Technology, property rights and organizational diversity in the software industry. Structural Change and Economic Dynamics 23 (2), 137–150.

Leiner, B. M., Cerf, V. G., Clarck, D. D., Kahn, R. E., Kleinrock, L., Lynch, D. C., Postel, J., Roberts, L. G., Wolff, S. S., 2001. The past and future history of the internet. Communication of the ACM 40 (2), 102–108.

Lessig, L., 1996. The zones of cyberspace. Stanford Law Review 48 (5), 1403–1411.

Lessig, L., 1999. Code and Other Laws of Cyberspace. Basic Books, New York.

Lessig, L., 2006. Code: version 2.0. Basic Books, New York.

MacKinnon, R., 2012. Consent of the Networked: The Worldwide Struggle for Internet Freedom. Basic Books, New York.

Marx, K., 1970. Il Capitale: Critica dell'Economia Politica. Newton Compton editori, Roma.

McManis, C. R., 2008. The proposed anti-counterfeiting trade agreement (acta): Two tales of a treaty. Houston Law Review 46 (4), 1235–1256.

Mitchell, W. J., 1995. City of Bits: Space, Place, and the Infobahn. MIT Press, Cambridge, MA.

Naidu, S., Hwang, S.-H., Bowles, S., 2010. Evolutionary bargaining with intentional idyosincratic play. Economic Letters forthcoming.

Parsons, T., 1963. On the concept of political power. Proceedings of the American Philosphical Society June, 232–262.

Post, D., 1995. Anarchy, state and the internet. Journal of Online Law (3).

Reidenberg, J. R., 1998. Lex informatica: The formulation of information policy rules through technology. Texas Law Review 76 (3), 553–593.

Shaw, A., 2008. The problem with the anti-countereiting trade agreement (and what to do about it). KEStudies 2.

Simon, H. A., 1951. A formal theory of the employment relationship. Econometrica 19 (3), 293–305.

Sterling, B., 2002. The Hacker Crackdown: Law and Disorder on the Electronic Frontier. Bantam Books, New York.

Strahilevitz, L. J., 2003. Charismatic code, social norms, and the emergence of cooperation on the file-swapping networks. Virginia Law Review 89 (3), 505–595.

von Hippel, E., 2005. Democratizing Innovation. MIT Press, London.

Weber, M., 1978. Economy and Society. University of California Press, Berkeley.

Wu, T., 2003a. Network neutrality, broadband discrimination. Journal on Telecommunication and High Technology Law 2, 141–176.

Wu, T., 2003b. When code isn't law. Virginia Law Review 89 (4), 679–751.

Wu, T., 2010. The Master Switch: The Rise and Fall of Information Empires. Random House Inc., New York.

Young, H. P., 1998. Individual Strategy and Social Structure: An Evolutionary Theory of Institutions. Princeton University Press, Princeton.

Zittrain, J., 2000. What publisher can teach the patient: Intellectual property and privacy in an era of trusted privication. Stanford Law Review 52 (5), 1201–1250.

Zittrain, J., 2003. Internet points of control. Boston College Law Review 44 (653).

Zittrain, J., 2006. The generative internet. Harvard Law Review 119 (7), 1974–2040.

Zittrain, J., 2008. The Future of the Internet and How to Stop It. Yale University Press, New Haven and London.