

Governance of Digital Content in the Era of Mass Participation¹

Tobias Regner, Javier A. Barria, Jeremy V. Pitt and Brendan Neville

Intelligent Systems & Networks group
Department of Electrical and Electronic Engineering
Imperial College London
Exhibition Road, SW7 2BT, London

{t.regner,j.barria,j.pitt,b.neville}@imperial.ac.uk

JEL classifications: L10, L82, L86, M21, P48, Z13

Key words: music industry, copyright, digital media, creative commons, multi-agent systems

Abstract:

The realm of digital media is undergoing fundamental changes as it moves from mass media to an era of mass participation. This emergence of content created by the masses requires to re-consider the conventional intellectual property rights framework. Free content and protected content co-exist (in the Dark/Light Web).

We propose an alternative environment for the governance of digital content. It incorporates psychological aspects into its economics framework. Multi-agent systems play an important role in order to create an infrastructure that makes the voluntary-based environment sustainable.

We propose a platform based on an open contracts design that encourages voluntary payments. Peer-based reputation and recommendation mechanisms as well as socio-emotive instruments facilitate norm-adherence in this online environment. They leverage the efficiency of alternative voluntary payment models based on fairness concerns and reciprocity.

The envisioned platform matches Dark Web content to consumers who value it highly, provides Dark Web content creators with a basic reward for their work and reduces the infringement of protected content in the Dark Web.

1. Introduction

Copyright law originated in the 18th century and the protection it provides is regarded as an essential cornerstone for a successful intellectual property rights environment. However, recently it has come under increasing scrutiny. Its original intention of promoting the “progress of science and of useful arts”

¹ This work was partially supported by the EU-funded project DANAE (IST-2004-507113).

by granting a temporary monopoly has been diluted significantly over the years. It has become an instrument of commercial interests that now offers over 100 years of protection, even retroactively. Moreover, the attempts to adjust copyright law to the digital age — the Digital Millennium Copyright Act in the USA and the EU Copyright Directive — essentially eliminated any fair use rights granted by the original copyright law.

While the interests of consumers and the content industry have by and large been fairly balanced by copyright law during the analogue era, the current legislation does not seem to take established consumption patterns of the past and the potential benefits of online distribution into account. Doubts remain whether copyright law is still an adequate governance system for intellectual property rights in the digital age and whether there are alternative ways to govern digital content, a more efficient eco-system for ideas than the one copyright law offers?

As the balance of copyright law shifted towards industry interests, an alternative rights scheme – the creative commons license – emerged that provides more flexibility to content owners and addresses the recent negligence of consumer rights. Its success² underlines the emergence and importance of content created by the masses. This development leads to a two-poled digital media environment in which the two domains – free and protected content or the Dark and the Light Web – do not compete, but co-exist and complement each other. (Regner et al. 2006) While content created by big media (e.g. famous, established artists) is highly commercialised and embraces protection by conventional copyright law, content created by the masses (e.g. amateurs and newcomer artists) prefers to be available under more flexible terms providing better exposure.

Since the pool of free Dark Web content is vast and access to it is free, it remains to be seen how content can be searched and found effectively and how its creators can at least be rewarded on a basic level to make their efforts worthwhile. Unless these two aspects are not solved with satisfaction, the Dark Web either becomes a big unorganised sprawl of data or it dries out because of a lack of incentives. Moreover, infringement of copyrighted content must be kept in check in the Dark Web.

Hence, the paper's aim is to analyse the emerging hybrid copyright world with free and protected content environments co-existing. We describe how the potential of the Dark Web – the free content environment – can be reached and what is necessary to make the Dark Web 1) search-efficient, 2) financially sustainable and 3) norm adhering.

Our approach is inter-disciplinary as socio-economic, technical as well as legal aspects are taken into account. It builds on the concept of an open agent society for digital content proposed by Pitt (2005). We suggest a platform for the provision of digital content that incorporates psychological aspects into its economics framework. Multi-agent based reputation and recommendation systems play an important role in order to create an infrastructure that makes voluntary payment-based models sustainable.

² The creative commons license is used by 45 million web sites and has been adapted to the legal systems of 29 countries as of December 2005.

The contribution of behavioural economics is a more realistic analysis of economic situations, one that takes the psychological underpinnings of economic behaviour better into account. Social preferences explain this behaviour as they "assume people are self-interested, but are also concerned about the payoffs of others." (Charness and Rabin 2002) This departure from pure self-interest is based on the results of a vast number of laboratory experiments conducted in recent years. (Camerer 2003, Fehr and Schmidt 2005) Also, Regner and Barria (2005) found a significant level of voluntary payments in the real-life environment of an online music store which is particularly relevant to the here analysed context.

Multi agent-based systems can complement this as intelligent information agents search, find and retrieve content with respect to users' individual preferences and provide the peer-based infrastructure to encourage pro-social behaviour and punish free riding.

From a legal perspective, this relates to the literature about legal rules and standards. (Burk 2005) Copyright law with all its exceptions (e.g. fair use) is essentially a case of a legal standard as it grants substantial discretion to consumers. Digital rights management (DRM) systems resemble legal rule making as they shift discretion to producers. The intended exceptions of original copyright law cannot be taken into account by rigid DRM designs. A softer, more flexible system would be preferable to govern intellectual property rights in a socially efficient way.

Such a system also focuses on conflict prevention. It aims to solve legal disagreements about rights infringement efficiently before they actually make it to the formal system of lawyers and courts. This would mean substantial cost savings to society since conflicts are prevented before the legal system comes into action.

Section 2 describes the changes in the digital media industry and its emerging two-poled structure. It also explains the peculiarities of content created by the masses for the masses and its implications. Section 3 is about the theoretical background. Section 4 describes the envisaged platform and Section 5 concludes.

2. Industry Snapshot

Two domains co-exist in the hybrid copyright world that is emerging. DRM-protected content is available in the Light Web, while the Dark Web contains free content created by the masses.

The Light Web is essentially the conventional offline retail model moved online. Established artists (and their label) benefit from the high level of content protection provided by the Light Web environment. This guarantees constant revenue streams as long as the artist remains popular. The Light Web also means exposure of the content is restricted (through distribution limits and higher prices than free content) which is ideally offset by the established reputation and sizeable fan base of famous artists.

However, protected content's limitations on the potential exposure would be a concern for newcomer artists. Instead, they appreciate the informational role of free content distribution, since it is their only way to find a big audience

without the conventional label promotion. It is also the domain for amateur content that is increasingly put online without any (immediate) commercial motivations.

The two domains, Light Web and Dark Web, do not compete against each other for domination of the digital content market. Instead, both provide unique advantages for certain types of artists and are therefore inhabited by different types of artists. While exposure appears to be essential for newcomers to start off, guaranteed revenues and therefore stronger content protection will likely affect the decision of established artists. Voluntary-based models will appeal to newcomers (and amateurs) in the Dark Web. DRM-based models are the likely choice of established artists in the Light Web. There is a flow or interaction between the two environments when newcomer artists of the Dark Web are discovered by the masses and eventually turn to a more commercial Light Web approach with a higher degree of protection and the guaranteed revenues that protected content provides. Complements-based³ or even voluntary-based models can be the choice of famous artists, if they return to the Dark Web.

At the core of both domains lies the question of how the content distribution and search will be organised. In particular, the Dark Web requires an intelligent search infrastructure so that its free content will be found and consumed by customers who appreciate it most.

While business actors known from the traditional media industry (artists, labels) still play a role in this changing environment, it is innovative service providers that take on more prominent roles. Sales platforms like the iTunes Music Store, P2P search platforms and social network platforms like Yahoo! Music add value by assisting consumers to find the music they want. Social network platforms aim to provide their users with added value, for instance by making personalised music recommendations.

Two strands of music recommendation services can be distinguished: collaborative filtering and acoustics-based approaches.

Collaborative methods to recommend content based on the aggregated user ratings of this content have been first developed in the 90s, see for instance Shardanand and Maes (1995). The collaborative filtering predictions for a user are based on the similarity between the preference profile of that user and those of other available users. Amazon's collaborative filtering-based recommendation system uses purchase-level data and the limitations of its prediction accuracy are well known. More advanced approaches include the Audioscrobbler service (www.last.fm/index.php) which uses consumption-level data for its collaborative filtering as it tracks the listening habits of users that send in their information.

Another example is the music recommendation service of MoodLogic (www.moodlogic.com) which is survey-based. Arguably the world's largest collection of human data points about music, it contains mood-related questionnaire responses like opinions and feelings about songs. These song profiles are then used for mood-related recommendations.

³ Business models where content is distributed freely and revenue is generated from complementary products and services like concerts, merchandising and licenses.

Acoustics-based approaches aim to determine the acoustic similarity between songs. Recommendations can then be made based on a favourite song just heard and its acoustic 'neighbours'. The Songsphere service (www.bmat.com) follows this approach as it only uses musical features automatically extracted from audio signals, such as instrumentation, rhythm and harmony. Similarity profiles and recommendations are generated based on this audio data, without the use of meta data or human ratings.

3. Theoretical Background

The potential of voluntary payment schemes relies on the tendency of customers to have social preferences. While other-regarding behaviour occurs at significant levels in lab experiments as well as in field studies, a voluntary-based system may need to be complemented by additional instruments to be evolutionary stable and sustainable. Therefore, a peer-based system of checks and balances can be useful as it can provide further soft and hard incentives to adhere to norms and behave pro-socially. It may also help to prevent copyright infringement at a very early stage.

The concept of social preferences goes back to the very beginning of modern economics — in fact, literally even beyond that. Even before the “Wealth of Nations” Adam Smith already stressed the importance of other-regarding behaviour in his “Theory of Moral Sentiments”.

Social preferences signify a more realistic analysis of economic situations, one that takes the psychological underpinnings of economic behaviour better into account. This is achieved by modifying the standard neoclassical theory, essentially making its assumptions better suited to the psychological realities of economics. Other-regarding behaviour for instance complements self-interest in a theory of social utility. (Camerer 2003) While individuals with social preferences are driven by self-interest, they also take the well-being and/or the intentions of other individuals into account. This can include fairness considerations, reciprocity and concerns for social welfare (instead of private profit), but also the impact of emotional aspects (guilt, shame, pride) on decisions can elicit social preferences.

Without a doubt economic motivation by self-interest accurately explains behaviour in many situations. Predictions are particularly fine the more competitive markets are. However, many economic transactions are not taking place in a competitive, impersonal market environment and the more social and personal the exchange is, the more other-regarding behaviour matters. (Smith 1998)

Numerous laboratory experiments successfully tested for other-regarding behaviour among subjects.⁴ These results can be explained by reciprocity, an intentions-based approach to theoretically model social preferences. Kind behaviour is returned, while nasty actions are answered with spiteful behaviour. Therefore, an open contract design that provides mutual

⁴ For instance, Fehr, Gächter, Kirchsteiger (1997) found increased effort levels when mutual opportunities to reciprocate have been provided in experimental settings in a labour context (first a firm sets a wage, then a worker decides on the effort level, while finally the firm may pay a bonus).

opportunities to reciprocate appears to be a useful construct, in particular when effort – or payment in the context of digital media – cannot be enforced. The basic design can be extended so that users can punish free riders (users who frequently failed to contribute/pay). The effect of the mere existence of a way to punish un-cooperative users has been found to be beneficial as compliance increases. (Fehr and Gächter 2002)

Such a punishment option can sustain high cooperation rates, but it also causes negative welfare effects when this option is abused by anger-dominated subjects and rounds of retaliation result. (Hopfensitz and Reuben 2005) Instead, triggering shame feelings by increasing self-awareness should only cause remorseful subjects to cooperate (more), without the negative side effects known from the punishment option. The sequencing of the emotive instruments therefore appears to be important. While a punishment option might be necessary „for stability“ and to keep users honest, it should follow the self-awareness instrument that appeals to subjects to correct their behaviour themselves.

Multi agent-based systems are applied to digital content and the information society in Pitt (2005) by forming an open agent society. The key idea here is that if intelligence is encapsulated in an agent, then the agents should be organised as a society, and then the rules of the society can be used to ensure ‘responsible’ behaviour. It is discussed how an agent society can be used to regulate behaviour in future information trading scenarios and stressed that it is essential that we support information trading agents with decentralised mechanisms for enforcing honest and reliable behaviour.

Neville and Pitt (2004) propose a computational socio-cognitive and economic framework (CSCEF) as such a mechanism for distributed regulation. This inter-disciplinary framework formalises social theories of trust, reputation, recommendation and learning from direct experience and integrates these socio-cognitive formalisms with the agents’ economic reasoning. This produces an agent whose behaviour in commercial transactions is influenced by its social interactions, whilst being motivated and constrained by its economic considerations.

In the legal context, an approach towards intellectual property rights as a legal standard instead of a legal rule is preferable. Under the flexible governance of a legal standard reasonable exceptions can be granted and trivial violations are not pointlessly prosecuted. In contrast, a legal rule is strictly set up ex ante and leaves no leeway for interpretation. The realm of copyright is a field where a standard should be preferred. (Burk 2005) This was recognised by copyright law in its original form as it allows a number of exceptions (e.g. fair use) that can be handled under ex post discretion. DRM systems based on written copyright law form a legal rule instead. They prevent any actions that are technically an infringement and thus fail to allow for the exceptions incorporated into copyright law and the “wobble room” at the discretion of society.

In that sense, such a governance system would focus on conflict prevention which can be understood as the attempt to solve legal disagreements between parties before they become real disputes in front of the court and require dispute resolution or even litigation. Such a focus to solve cases

before they actually make it to the formal legal system would mean substantial cost savings to society since the full legal apparatus does not even get started.

4. The Suggested Platform

As explained before an open contracts design provides mutual opportunities to reciprocate and can encourage users to make voluntary payments. When content owners decide to make their content freely available through the Dark Web, they benefit from increased exposure and they may also expect a basic return of voluntary payments for their work that can cover their opportunity costs or may even exceed them. However, an open contracts design on its own is unlikely to be successful, because content may not be found in the vast Dark Web and voluntary payments may not reach a sufficient or stable level. Also abuse of copyrighted material may occur in the Dark Web. Further instruments are necessary to make the platform 1) search-efficient, 2) financially sustainable and 3) norm adhering.

In our suggested platform a multi agent-based infrastructure complements the open contracts design. It acts as the users' portal to the online content community as it facilitates the finding of interesting new content, the fair remuneration of consumed content and the norm adherence within the community.

While content is obtained directly from the creator (or his service provider) in the Light Web, it is shared amongst peers in the Dark Web. Recommendations about content are exchanged among peers to facilitate the matching of preferences. Moreover, recommendations about peer behaviour are shared among peers and content creators to build up trust ratings of peers. This information can be used to reciprocate positively against fair and negatively against unfair peers.

4.1. Search

The search efficiency for new good music is increased through automatically generated content recommendations within the community. As peers share their music preferences the taste matching of music amongst peers is improved. In contrast to conventional collaborative filtering which is based on purchase decisions, the content search can analyse the preference profiles of peers at a deeper level. Consumption data can be used instead of purchase data. The Shared Preferences Model⁵ (SPM) could use the ratings given to songs. However, it may be too cumbersome for many users to tag songs with their ratings and update them. Instead, using the play count of songs (and related data as frequency and time of the day etc.) promises to supply better data and a superior image of peers' music taste. Sharing the preferences with the P2P network should result in valuable recommendations of appreciated music, provided by a personal DJ.

4.2. Payment

⁵ For further description see Witkowski et al. (2003).

Similar to when the agent handles making content recommendations there may be a number of different metrics for estimating a suitable voluntary payment for the peer to make. These variables could include the:

- Number of times the content has been played
- Peer's explicit rating
- Content's relative position against this month's other new music
- Content owner's requested terms (expected payment)
- Content's genre/media category
- Peer's budget

Some of these variables can be specified by the peer. For instance, a certain (monthly) budget can be set that should be spent on content. Peers may also have different settings for voluntary payments depending on the genre. They may not wish to contribute to pop music, but feel more generously towards alternative music.

Content owners may suggest an expected payment for their content. It may be that the consumer doesn't view/download content of this type or simply ignores the suggestion (see consequences of norm violation in the next subsection). The agent could ignore suggestions that are outside of what is considered acceptable terms by either the online community as a whole or its own peer. Content owners can also set a very low expected payment to maximise their exposure. This would be the case of the Dark Web and its participatory media where the content owner may primarily be a consumer himself.

The agent will formulate suggested voluntary payments for each digital item on behalf of the user. At specified dates (e.g. monthly) or after a certain number of payments are ready to be made the agent will prepare a list of suggested payments to the peer for approval. Before any payments are authorised, the peer would be given the opportunity to view and if necessary alter the payments from those suggested by its agent. This preview screen would also allow the peer to query why the agent suggests any given payment. The peer can either accept or adjust their agents' suggested payments. Adjustments of the suggested payments will get fed back to customise the voluntary payment estimation algorithm in order for the agent to make more suitable suggestions in the future. If there are any payments authorised, then the agent performs the necessary micro-payments and stores the peer's receipts and newly acquired rights.

4.3. Norm Violation and Adherence

In determining whether a peer has violated a norm of the society, the sharing peer's agent is looking for a discrepancy between the peer agent's actions and the peer agent's statements. Such a discrepancy occurs for instance, if the peer states he enjoyed the content or if the peer is sharing the content without having made a payment. However, the peer may attempt to cover up the norm violation by either failing to rate the content or by making an incorrect negative recommendation. While this strategy protects the peer's trust rating (he avoids being labelled as a free rider, which would cause

negative reciprocity), it misinforms the shared preference model (biased recommendations would be made due to the wrong preferences provided). Therefore, agents who lie to avoid paying for content lose the added value of a well functioning recommendation system. A peer will never meet friends (other peers with similar tastes), if he pretends to dislike everything. He will only be ranked highly with those peers who also claim not to like anything. However, they won't be able to share content because they cannot communicate any likes or dislikes in order to do so.

Explicit use cases can be defined for norm violation, norm adherence and the update of trust ratings:

Following the sharing of content with a peer, the peer agent made:

No recommendation and no payment

In this case the sharing agent assumes after a period of time that the content was not of interest to its peer and automatically updates its Shared Preference Model (SPM) accordingly. In the future the peer will be offered less content of this type.

No recommendation and a payment

Upon receiving the notification of payment the sharing peer assumes the content was of interest to its peer and automatically updates its SPM accordingly. In the future the peer will be offered more content of this type from the sharing agent (but not the community as a whole as this would require a recommendation). The trust rating of this peer is improved so it could receive preferential treatment in the future.

Negative recommendation and no payment

The sharing peer updates its SPM so that the peer is offered less content of this type in the future.

Negative recommendation and a payment

The sharing peer updates its SPM so that the peer is offered less content of this type in the future. The trust rating of this peer is improved so it could receive preferential treatment in the future.

Positive recommendation and no payment

The peer has been identified as a free rider; this is an example of open norm violation. The sharing peer will update its trust rating of the peer to ensure that a retaliatory sanction is taken against the peer. For instance, in the future the agent will not be offered the certain content even though it is being shared with other peers.

Positive recommendation and a payment

The peer will update its SPM so in the future the peer will get more content of this type. The trust rating of this peer is improved so it could receive preferential treatment in the future.

Shared the content but made no payment

The peer has been identified as a free rider; this is an example of open norm violation. The sharing peer will update its trust rating of the peer to ensure that a retaliatory sanction is taken against the peer. For instance, in the future the agent will not be offered certain content even though it is being shared with other peers.

Potentially, the performance of this open norm-based community platform is threatened by malicious or naive agents. These undesirable agents may spread viral infections, masquerade as content producers but sell non-working files, execute denial of service attacks by constantly making random dataset comparison requests or unfairly rate other peers via the recommendation system. However, it is shown⁶ that by enabling agents to model their trust in a peer and make recommendations (which will lead to the formation of reputation beliefs) agents can isolate such harmful agents from the community, thus protecting themselves and the system performance as a whole.

As mutual opportunities to reciprocate increase the efficiency of an open contracts design, additional options to reciprocate negatively (punishment of free riders) or positively (rewarding frequent supporters) should increase the cooperation rate, e.g. the voluntary payments made. Hence, the platform's functionality should be extended for norm-adhering peers, while it should be restricted for norm-violators.

Once content is used in an inappropriate way agents interact with peers. Peers who regularly consume content without paying for it accordingly⁷ can be identified and punished by their peers. Their download speed can be reduced as well as the actual access to content. Harsh free riders may even be temporarily banned from the community. On the other hand, bonus or reward schemes can be implemented for consumers who frequently make voluntary payments or consumers who successfully recommend the creator. These benefits could be (backstage) concert tickets, limited edition T-shirts, autographs or related merchandise.

Moreover, emotional aspects like shame and embarrassment can be considered to make peers more aware of their 'negative' actions. In addition, forgiveness algorithms⁸ that take respective circumstances (past behaviour patterns, for instance) into account give free riders a chance to re-consider their uncooperative action, before steps to punish the user may be taken. Combined with P2P-based reputation and recommendation systems a sustainable environment can be created where abuse of free content is gently pointed out to the user and infringement of copyrighted material is limited. These instruments will leverage the efficiency of alternative voluntary payment models based on fairness concerns and reciprocity. Positive incentives are created to make peers conform to the rules of the system as pro-social behaviour is encouraged.

⁶ See the P2P trust and reputation framework developed in Neville and Pitt (2004).

⁷ Voluntary payment pleas may have been frequently disregarded or copyrighted material may have been downloaded and frequently consumed.

⁸ See Vasalou et al. (2006) for a forgiveness model.

Monitoring systems based on statistics, variation or historic patterns of interaction can also help to track down potential disputes about copyrighted material and assist to negotiate a way back to where both sides benefit. Conflict prevention can also be decentralised when peers of a 'community' act themselves against unsocial behaviour by other peers.

5. Conclusions

The emergence of participatory media⁹ – which contrasts the incumbent mass media approach – stresses the importance of an open and consumer-oriented solution to the way digital content is governed.

A flexible platform – suitable for free and protected content – becomes of increasing importance in the realm of digital media where a hybrid copyright world is about to develop. The emergence of content created by the masses and subsequently two co-existing and complementary environments for digital content require that content activities along the value chain also need to cater for "free" content created by the masses besides conventional protected content. The development of efficient tools and instruments in the content finding phase (reputation and recommendation systems, collaborative filtering) as well as voluntary-based business models should be of benefit.

The Dark Web platform we propose is based on an open contracts design that encourages voluntary payments. Peer-based reputation and recommendation mechanisms as well as socio-emotive instruments facilitate norm-adherence and pro-social behaviour in this online environment. They will leverage the efficiency of alternative voluntary payment models based on fairness concerns and reciprocity. The envisioned platform aims to match Dark Web content to consumers who value it highly, provide Dark Web content creators with a basic reward for their work and reduce the infringement of protected content in the Dark Web.

Such an environment helps to reap the potential benefits of an information society. The approach does not tackle the challenges faced by the digital media industry merely by battling technological progress and locking away content as is the case with most DRM systems and "trusted computing". Instead, a socially efficient solution is the goal. The proposed platform can be seen as a better way to align corporate and consumer interests in digital media. Its components describe what current systems may need to adopt to reach this goal.

6. Bibliography

Burk, D. L. (2005), "Legal and Technical Standards in Digital Rights Management Technology", *74 Fordham L. Rev.* 537 (2005)

Camerer, C. (2003), "Behavioral Game Theory: Experiments in Strategic Interaction", Princeton University Press, Princeton

⁹ See Kluth (2006) for more about the coming era of mass participation.

Charness, G., and M. Rabin (2002), "Understanding Social Preferences with Simple Tests", *Quarterly Journal of Economics*, 117, 817-869

Fehr, E. and S. Gächter (2000), "Cooperation and Punishment in Public Goods Experiments", *American Economic Review*, 90, 980-994

Fehr, E., S. Gächter, and G. Kirchsteiger (1997), "Reciprocity as a Contract Enforcement Device: Experimental Evidence", *Econometrica*, 65, 833-860

Fehr, E. and K. Schmidt (2005), "The Economics of Fairness, Reciprocity and Altruism - Experimental Evidence and New Theories", *University of Munich discussion paper*, June 2005

Hopfensitz, A. and E. Reuben (2005), "The Importance of Emotions for the Effectiveness of Social Punishment", *Tinbergen Institute Discussion Paper*

Kluth, A. (2006), "Among The Audience", *The Economist*, April 22nd 2006

Neville, B. and J. V. Pitt. (2004), "A Simulation Study of Social Agents in Agent Mediated E-Commerce. " In: Falcone et al., *Proc. of the AAMAS-2004 Workshop on Trust in Agent Societies*, pages 83-91.

Pitt, J. V. (2005), "The Open Agent Society as a Platform for the User-friendly Information Society", *AI & Soc* (2005), 19: 123–158

Regner, T. and J. A. Barria (2005), "Magnatune: Variable Pricing for Music," mimeo, online available at:
http://papers.ssrn.com/sol3/papers.cfm?abstract_id=721596

Regner, T., J. A. Barria, J. V. Pitt and B. Neville (2006), "Emerging Business Models in Digital Media", mimeo, online available at:
http://papers.ssrn.com/sol3/papers.cfm?abstract_id=892792

Shardanand, U. and Maes, P. (1995), "Social Information Filtering: Algorithms for Automating 'Word of Mouth'", *Proceedings of the Conference on Human Factors in Computing (CHI95)*, 210-217, Denver, ACM Press

Smith, V. L. (1998), "The Two Faces of Adam Smith", *Southern Economic Journal*, 65, 2-19.

Vasalou, A., J. V. Pitt and G. Piolle (2006), "From Theory to Practice: Offering Forgiveness as a Way to Repair Online Conflicts in CMC", *Proceedings of the iTrust Conference*, Pisa, Italy, 2006

Witkowski, M., B.Neville and J. Pitt (2003), "Agent Mediated Retailing in the Connected Local Community", *Interacting with Computers*, 15(1):5-32