To appear in *Proceedings of Second International Workshop on Cooperative Information Agents (CIA-98)*, Paris, July 4--7, 1998.

Dynamics of an Information-Filtering Economy

Jeffrey O. Kephart, James E. Hanson, David W. Levine, Benjamin N. Grosof, Jakka Sairamesh, Richard B. Segal, and Steve R. White

IBM Thomas J. Watson Research Center
P.O. Box 704, Yorktown Heights, NY 10598, USA
{kephart,hanson,dwl,grosof,jramesh,rsegal,srwhite}@watson.ibm.com

Abstract. Our overall goal is to characterize and understand the dynamic behavior of information economies: very large open economies of automated information agents that are likely to come into existence on the Internet. Here we model a simple information-filtering economy in which broker agents sell selected articles to a subscribed set of consumers. Analysis and simulation of this model reveal the existence of both desirable and undesirable phenomena, and give some insight into their nature and the conditions under which they occur. In particular, efficient self-organization of the broker population into specialized niches can occur when communication and processing costs are neither too high nor too low, but endless price wars can undermine this desirable state of affairs.

1 Introduction

Today, we are witnessing the first steps in the evolution of the Internet towards an open, free-market *information economy* of automated agents buying and selling a rich variety of information goods and services[2, 4, 6, 16, 19, 22]. We envision the Internet some years hence as a seething milieu in which billions of economically-motivated agents find and process information and disseminate it to humans and, increasingly, to other agents. Over time, agents will progress naturally from being mere facilitators of electronic commerce transactions to being financial decision-makers, at first directly controlled by humans and later with increasing autonomy and responsibility. Ultimately, inter-agent economic transactions may become an inseparable and perhaps dominant portion of the world economy.

The evolution of the Internet into an information economy seems as desirable as it does inevitable. After all, economic mechanisms are arguably the best known way to adjudicate and satisfy the conflicting needs of billions of *human* agents. It is tempting to wave the Invisible Hand and assume that the same mechanisms will automatically carry over to *software* agents. However, automated agents are *not* people! They make decisions and act on them at a vastly greater speed, they are immeasurably less sophisticated, less flexible, less able to learn, and notoriously lacking in "common sense". How might these differences affect the efficiency and stability of future information economies?

Previous research in automated economies is equivocal. Under certain assumptions, large systems of interacting self-motivated software agents can be susceptible to the emergence of wild, unpredictable, disastrous collective behavior[13, 14]. On the other hand, a large body of work on market mechanisms in distributed multi-agent environments suggests that efficient resource allocation or other desirable global properties may emerge from the collective interactions of individual agents[1, 3, 5, 8, 10, 11, 15, 17, 21].

Our goal is to understand the dynamic, emergent behaviors — both good and bad — of information economies from an *agent's*-eye view, and from this to formulate basic design principles that will foster efficient electronic commerce. We pursue this goal by combining analysis and simulation of information economies with concurrent development of an information economy prototype.

In this paper, we focus on a simple model of an information filtering economy, such as might be embedded in a larger information economy. The model is inspired by information dissemination services that can be found on the Internet today, and sets them in an economic context. After introducing the model in section 2, we analyze and simulate its dynamical behavior in section 3, illustrating as we go the promise and the pitfalls inherent in this and similar economies. We conclude with a brief summary of our findings in section 4.

2 Model of the news filtering economy

Fig. 1 represents our information filtering model economy, consisting of a *source* agent that publishes news articles, C *consumer* agents that want to buy articles they are interested in, B *broker* agents that buy selected articles from the source and resell them to consumers, and a *market* infrastructure that provides communication and computation services to all agents. Each agent's internal parameters (defined below) are printed inside its ellipse. Solid lines represent the propagation of a sample article through broker 1. Broken lines indicate payment, and are labeled with symbols (explained below) for the amount paid.

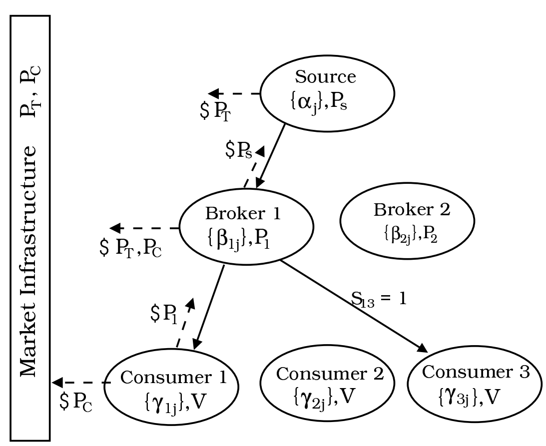

Fig. 1. Part of an idealized news filtering economy. Only a subset of agents is shown. See text for interpretation of symbols.

The source agent publishes one article at each time step t. It classifies articles according to its own internal categorization scheme, assigning each a *category index* j. The nature of the categories, and the number J of them, do not change. We represent this (hidden) classification scheme by a random process in which an article is assigned category j with fixed probability α_j. The set of all α_j is the source's *category prevalence* vector α. Each article is labeled with its category index and offered for sale to all brokers at a fixed price P_S. For each article sold to each broker, the source pays a fixed *transport cost* P_T.

Upon receiving an offer, each broker b decides whether or not to buy the article using its own evaluation method, which may be *uncorrelated* with the source's categorization scheme. For each evaluation that it makes, the broker pays the system a fixed *computation cost* P_C. The broker's evaluation method is approximated by a random process parametrized by its *interest vector* β_b: it buys an article labeled (by the source) with category j with probability β_{bj}. When broker b purchases an article, it immediately sends it to a set of *subscribing* consumers, paying tranportation cost P_T for each. Subscribers examine the article, and pay the broker P_b if they want the right to use ("consume") it. The broker's internal parameters β_b and P_b are under its direct control.

Subscriptions are represented by a *subscription matrix* S, where $S_{bc} = 1$ if consumer c subscribes to broker b, and $S_{bc} = 0$ if not. Subscriptions are maintained only with the consent of both parties, and may be cancelled by either. For example, a broker b might reject a consumer c if the cost of sending articles exceeds the expected payment from c, or c might reject b if the cost of sifting through lots of junk outweighs the benefit of receiving the rare interesting article. The bilateral nature of the agreement is represented by setting $S_{bc} = \sigma_{bc}^{(b)}\sigma_{bc}^{(c)}$, where $\sigma_{bc}^{(b)} = 1$ if broker b wants consumer c as a subscriber and 0 if not; analogously, $\sigma_{bc}^{(c)}$ represents consumer c's wishes.

When a consumer receives one or more copies of an article from brokers to which it subscribes, it pays the computation cost P_C to determine whether it is *interested* in the article, then decides whether (and from whom) to buy it. Like the brokers, the consumers' evaluation function is approximated by a stochastic process parametrized by an interest vector γ_c: consumer c will be interested in an article labeled with category j with fixed probability γ_{cj}. The consumer then determines whether the *anticipated value* V for interesting articles warrants paying the price P_{b^*} demanded by the chosen broker b^*, and if so purchases the usage rights to the article.

An alternative formulation replaces the consumer's computational cost P_C with a negative value or cost P_J incurred when "junk" is received. The transformation $V \to V + P_J$, $P_C \to P_J$ renders these two views equivalent.

3 Behavior of the news filtering economy

In this section we illustrate desirable and undesirable phenomena that can occur in our news filtering economy. First, we define the *state* of the system, from which any desired aspect of behavior can be derived. Then, we derive the state and behavior of simple systems with a few well-informed brokers and an infinite number of consumers. Finally, we simulate a system of many brokers and consumers with limited knowledge of the system state, and show that it can self-organize into a configuration that is beneficial to brokers and consumers alike.

We define the *state* of the system at time t, $\mathcal{Z}(t)$, as the collection of broker prices P_b, broker interest vectors β_b, and the subscription topology matrix S. Our goal is to understand the evolution of $\mathcal{Z}(t)$, given (i) an initial configuration $\mathcal{Z}(0)$; (ii) the values of the various extrinsic (possibly time-varying) variables: the category prevalences α_j, the costs P_S, P_T, and P_C, the consumer value V, and the consumer interest vectors γ_c; and (iii) the algorithms used by each agent to dynamically modify those variables over which it has control, including specification of a) the state information accessible to the agent and b) the times at which the modifications are made.

Any desired individual or aggregate aspect of behavior can be extracted from the history of $\mathcal{Z}(t)$ and the extrinsic variables. Two particularly important quantities are the expected utility per article for consumers and brokers. It can be shown that the expected utility per article U_c for consumer c is given by:

$$U_c = \sum_{b=1}^{B} \sum_{j=1}^{J} \alpha_j S_{bc} \beta_{bj} \left(\prod_{b'=1}^{B} (1 - S_{b'c} \beta_{b'j} \Theta(P_b - P_{b'})) \right) [(V - P_b)\gamma_{cj} - P_C] \quad (1)$$

where $\Theta(x)$ is the step function: $\Theta(x) = 1$ for $x > 0$, and 0 otherwise. The product term in large parentheses is the probability that an article in category j is not offered by any broker for a price less than P_b. The term in square brackets is the expected value of an article in category j: it always costs the consumer P_C to process it, regardless of its worth, and with probability γ_c consumer c will pay P_b to receive information worth V.

The appropriate utility function for the broker is its expected profit per article, given by:

$$W_b = \sum_{c=1}^{C} \sum_{j=1}^{J} \alpha_j S_{bc} \beta_{bj} \left[P_b \gamma_{cj} \left(\prod_{b'=1}^{B} (1 - S_{b'c} \beta_{b'j} \Theta(P_b - P_{b'})) \right) - P_T \right] \quad (2)$$
$$- \sum_{j=1}^{J} \alpha_j (\beta_{bj} P_S + P_C)$$

3.1 Single broker case

Given that effective monopolies can occur even in multi-broker systems, it is useful and instructive to establish a few simple results for systems with a single broker. First, suppose the broker offers a single category, and that the number of consumers is arbitrarily large, i.e. $B = 1$, $J = 1$, and $C \to \infty$. The broker tries to maximize its utility by choosing a set of preferred consumers (those c for which $\sigma_c^{(b)} = 1$), setting a price P, and setting its interest level β. The consumers try to maximize their utility simply by declaring whether they wish to subscribe to the broker (in which case $\sigma_c^{(c)} = 1$).

Analysis of Eqs. 1 and 2 shows that, for a wide range of parameters, the equilibrium state is $\{\beta = 1, P = P^*, S_c = \Theta(P^*\gamma_c - P_T)\Theta((V - P^*)\gamma_c - P_C)$. The subscription matrix element S_c, which defines whether or not consumer c has subscribed to the broker, is the product of two step functions, which can be understood intuitively as follows. The first step function, $\Theta(P^*\gamma_c - P_T)$, represents the veto power of the broker: it only wishes to serve consumers that are interested enough (and the price is high enough) so that the expected revenue from an article will exceed the cost of sending it. The second step function, $\Theta((V - P^*)\gamma_c - P_C)$, represents the consumer's veto power: it only wishes to subscribe to the broker if it is interested enough (and the price is low enough) so that the expected net benefit of receiving an article exceeds the cost of processing it. The monopolistic equilibrium price P^* is constrained by the step functions and the restriction $\gamma_c < 1$ to be in the range $P_T < P^* < V - P_C$; its exact value depends in detail upon P_C, P_T, V, and the distribution of consumer interest probabilities $\Gamma(\gamma_c)$ in the population. For example, when Γ is a uniform distribution, P^* is the solution to a cubic equation involving P_C, P_T, and V.

Now suppose that the number of categories J offered by the broker is arbitrary [7]. Assuming that the category prevalences are all equal ($\alpha_j = 1/J$) and the distribution of consumer interests within a given category is given by Γ, one can derive analytic expressions for the monopolistic equilibrium price as a function of the number of categories. Substituting this function $P^*(J)$ back into Eq. 2, one can derive the broker's optimal utility as a function of the number of categories, and then from this the optimal number of categories. Illustrative results for two very different distributions Γ are shown in Fig. 2.

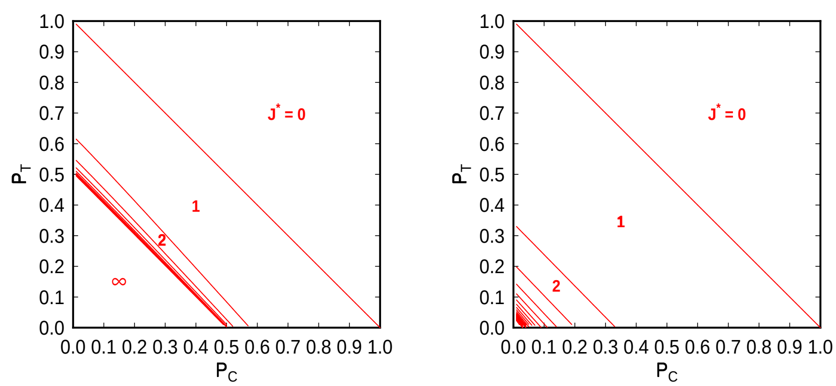

a) Uniform distribution b) All-or-nothing distribution, $\nu = 0.03$.

Fig. 2. The optimal number of categories J^* for the broker to offer as a function of P_C and P_T (with $V = 1$). a) Γ is a uniform distribution in the interval (0,1), so $\nu = 0.5$. b) Γ is an all-or-nothing distribution: $\gamma_j = 1$ with probability $\nu = 0.03$, else $\gamma_j = 0$.

For a wide class of distributions, three behavioral regimes are observed. When the combined cost of transport and processing is sufficiently high ($P_C + P_T > V$), the optimal number of categories J^* is 0. In this "dead" regime, an article costs more to send and process than it is worth, even if the consumer is guaranteed to be interested in it, and so no articles will be bought or sold. At the other extreme, when the costs are sufficiently low ($P_C + P_T < \nu V$, where ν is the mean of the distribution Γ), the broker is motivated to offer *all* categories that exist ($J^* \to \infty$). In real information filtering applications, one expects ν to be quite small, since each consumer regards most information as junk. It is useful to think of $J^* \to \infty$ as a (presumably tiny) **spam regime**, in which it costs so little to send information, and the financial impact on a consumer of receiving junk is so minimal, that it makes economic sense to send all articles to all consumers. In between these two regimes, the optimal number of categories is finite.[1]

3.2 Two broker case: price competition and warfare

To explore some effects of price competition, we begin by considering a simple two-broker system with a single information category. We assume that brokers and consumers are fully knowledgeable about the state of the system (in particular, they know the prices and interest vectors of all of the brokers). Furthermore, they instantly adjust their desired subscription vectors $\sigma_{bc}^{(b)}$ and $\sigma_{bc}^{(c)}$ to maximize their utility given the current set of prices and interest vectors. This last assumption removes the degrees of freedom associated with the subscription matrix by expressing it in terms of the other state parameters (prices and interest vectors) — a tremendous simplification.

We assume that the brokers update their prices asynchronously. One plausible strategy is for a broker to set its price to the value that maximizes its profit, assuming all other prices remain fixed. Such an update strategy is guaranteed to produce the optimal profit up until the moment when the next broker updates *its* parameters. We call such a strategy "myopically optimal", or **myoptimal**.

A useful construct for understanding the resulting dynamics is the *profit landscape*. We define a broker's profit landscape as its utility (given in Eq. 2) as a function of the prices offered by all brokers in the system, itself included. A contour map of the profit landscape for broker 1 in a system with $P_C = P_T = 0.3$, $V = 1$, and a uniform distribution of consumer interests is shown in Fig. 3a.

The landscape shown in Fig. 3a has two distinct humps. The "cheap" hump on the right corresponds to the case where broker 1 is cheaper ($p_1 < p_2$). Here, its profit is completely independent of p_2, and it can charge the monopolistic price derived previously. The "expensive" hump on the left corresponds to the case in which broker 1 is more expensive than broker 2, but still able to find customers. This comes about when broker 2 is charging so little that it cannot afford to keep marginal customers — i.e., customers with low interest levels γ_c — as subscribers. (Recall from the discussion of the single-broker case that a broker

[1] Closed-form expressions for the finite $J^* > 1$ contours appear impossible; these contours have been computed numerically in Fig. 2.

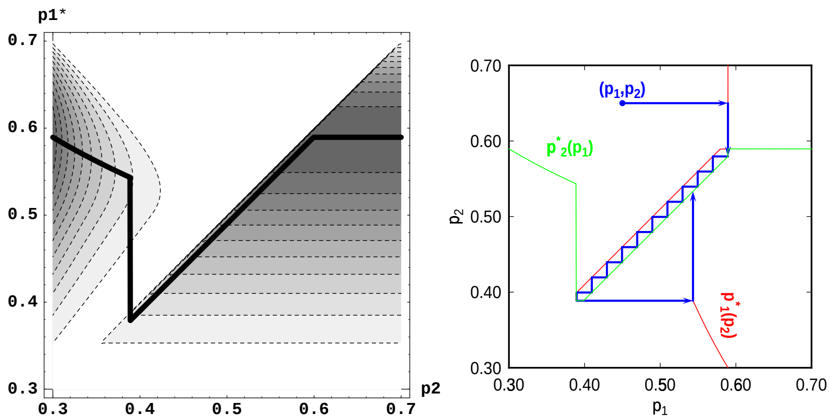

a) Landscape contours b) Price war construction

Fig. 3. a) Contour map of profit landscape for broker 1 for $P_T = P_C = 0.3, V = 1$, with overlaid optimal price function $p_1^*(p_2)$. Profit is higher in dark regions. b) Graphical construction of price-war time series, using functions $p_1^*(p_2)$ and $p_2^*(p_1)$.

will veto a subscription from an insufficiently interested prospective customer.) For these marginal customers, the only alternative is to subscribe to broker 1. In other words, broker 2's rejects constitute broker 1's market.

If broker 1 is myoptimal, it can derive from its profit landscape a function $p_1^*(p_2)$ that gives the value of p_1 that maximizes the profit when broker 2 charges price p_2. This function is represented as a heavy solid line in Fig. 3a. For $0.3 < p_2 < 0.389$, $p_1^*(p_2)$ is given by the solution to a cubic equation involving cube roots of square roots of p_2; in this region it looks fairly linear. The "vertical" segment at $p_2 = 0.389$ is a discontinuity as the optimal price jumps from the "expensive" hump to the "cheap" hump. When $0.389 < p_2 < 0.590$, $p_1^* = p_2 - \epsilon$, where ϵ is a price quantum — the minimal amount by which one price can exceed another. For $0.590 < p_2 < 0.7$, $p_1^* = 0.590$, the monopolistic price.

If broker 2 also uses a myoptimal strategy, then by symmetry its landscape and price-setting function $p_2^*(p_1)$ are identical under an interchange of p_1 and p_2. Then the evolution of both p_1 and p_2 can be obtained simply by alternate application of the two price optimization functions: broker 1 sets its price $p_1(t+1) = p_1^*(p_2(t))$, then broker 2 sets its price $p_2(t+2) = p_2^*(p_1(t+1))$, and so forth. The time series may be traced graphically on a plot of both $p_1^*(p_2)$ and $p_2^*(p_1)$ together, as shown in Fig. 3b. Assume *any* initial price vector (p_1, p_2), and suppose broker 1 is the first to move. Then the graphical construction starts by holding p_2 constant while moving horizontally to the curve for $p_1^*(p_2)$. Then, p_1 is held constant while moving vertically to the curve $p_2^*(p_1)$. Alternate horizontal moves to $p_1^*(p_2)$ and vertical moves to $p_2^*(p_1)$ always lead to a **price war** during which the brokers successively undercut each other, corresponding to zig-zagging between the diagonal segments of the curves. Eventually, the price gets driven down to 0.389, at which point the undercut broker (say broker 1) opts out of the price war, switching to the expensive hump in its profit landscape by setting $p_1^*(0.389) = 0.543$. This breaks the price war, but unfortunately it triggers a new one. The brokers are caught in a never-ending, disastrous limit cycle of price wars punctuated by abrupt resets, during which their time-averaged utility is half what they expected, and less than half of the monopolistic value.

3.3 General myoptimal case; discussion

Generalizing to an arbitrary number of brokers and categories, and permitting each broker to myoptimally update both its price *and* its interest vector, we observe more complex analogs of price wars, in which both prices and interest vectors are drawn into limit cycles. In the spam regime, the system tends to behave very wildly. When J^* is finite, the interest vectors can display some metastability, but price wars can develop even among brokers with different interest vectors (if they overlap sufficiently).

Price wars are even a problem when $J^* = 1$. Consider a system in this regime with n brokers and n categories. Such a system can accommodate each broker's wish to be a monopolist in a single category. If all categories are preferred equally by the users, each broker will ultimately specialize in a single unique category (even when J^* is somewhat more than 1) [7]. However, if the consumer population slightly favors one category, a system of niche monopolists is unstable, because each broker will cut its price in an effort to own the favored category. Simulations reveal that slightly less favored categories tend to be available much less often than the consumer population would like. Consequently, the total consumer utility is often reduced during a price war, despite the low prices [12].

Intuitively, any sort of economy consisting of myoptimal agents is likely to be plagued with price-war limit cycles. Geometrically, this behavior can be traced to the multi-peaked, discontinuous topology of the profit landscape, which in turn arises from the consumers' preference for the cheapest brokers. Of course, this does not imply that agent economies are doomed to failure. The assumption of myoptimality is unrealistic in several respects, and the fact that price wars occur relatively infrequently in human economies offers some hope.

The economics literature describes several possible mitigating effects that may explain why price wars are not rampant in human economies [18]. Expressed in terms of our model, these include explicit collusion among brokers, or tacit collusion — using foresight to avoid triggering price wars. Other factors thought to hinder price wars include frictional effects (consumers may find it too costly or bothersome to shop around, and brokers may find it costly to update prices or change products too often) or spatial or informational differentiation (i.e. different consumers might value the same good differently, depending on their physical location or knowledge).

These mitigating factors are likely to be weaker in agent-based economies than they are in human economies. Explicit collusion might require fairly sophisticated languages and protocols (and might be declared illegal!) In large decentralized systems, efforts to employ foresight may be hampered by imperfect knowledge of the system state and the strategies of the other agents. Even if these are known perfectly, it may be computationally infeasible to predict the future. [2] Consumer inertia may be greatly reduced when agents rather than people are doing the shopping, and price updates may be cheaper to compute and advertise. Localization effects are likely to be much smaller for information goods and services than they are for carrots and carwashes. Given these considerations, it is very possible that real agent economies will experience price wars much more frequently than do human economies.

3.4 Limited competitive knowledge; niches and prices

In order to understand the behavior of more realistic economies in which brokers and consumers are less informed about the system state, and less immediately responsive to environmental changes, we have built an agent simulation environment that allows us to experiment with a wide range of utility maximization strategies for consumers and brokers. In particular, we do not assume that brokers know each other's prices (or interest vectors), nor that the brokers know the consumers' interest vectors. Instead, they must make do with historical data based on their own parameters (e.g. prices and interest vectors) and experience (e.g. consumer demand and profits).

Here we describe a simulation run involving 10,000 consumers, 500 brokers, and 100 categories. In this run, a typical consumer was completely interested ($\gamma_{cj} = 1$) in roughly 3 out of the 100 categories and completely uninterested ($\gamma_{cj} = 0$) in all other categories. The computation and transport costs were $P_C = 30$ and $P_T = 30$; the information value was $V = 150$. Under these circumstances, $J^* = 1$. To set prices, the brokers use an extremely simple-minded strategy: they randomly shift their price by a small amount up or down. If (after a suitable period of time), the broker finds that its profit per unit time has increased, it keeps moving the price in the same direction; otherwise it reverses direction. We call this a *derivative-following* algorithm. Brokers adjust their interest vectors in a similar way, increasing or decreasing β_j by an amount depending upon the profit or loss recently experienced when selling an article in category j. The brokers and consumers estimate the utility of adding or cancelling subscriptions using a matchmaker that periodically gauges agents' interests by issuing questionnaires about a given set of articles.

Each consumer initially establishes a subscription with a single randomly chosen broker. From time to time, agents make asynchronous, independent decisions about adjusting prices, interest vectors, or subscriptions. Figure 4 shows the distribution of consumers' utilities at three different moments in their evolution. Starting from a state in which virtually all of the consumers have negative utilities, the economy adapts itself such that most consumers have positive utilities, and at worst a few have zero.

[2] In a simulation of just 4 brokers, 4 categories, and 10,000 consumers, computing a single *myoptimal* decision requires a few hours on a high-end workstation. Other researchers have shown that agents *can* learn to predict one another's actions, but the experiments have involved societies of fewer than ten agents [9, 20].

Of the original 500 brokers, only 122 remain active at time $t = 200,000$ (i.e., after 200,000 articles have flowed through the system). The others have either gone broke or are not participating in the market because they are buying no articles from the news source. Many brokers die out within the first few thousand time steps, while others thrive at first but then suddenly start losing to competition. *All* of the 122 that remain at $t = 200,000$ have chosen just a single category in which to specialize, and among them cover 86 distinct categories.

At $t = 0$, almost all of the consumers have negative utility because the random initial conditions typically do not give them a large enough ratio of interesting articles to junk. Rapid improvements are soon made as the agents mutually sort out the subscription topology. During this phase, many consumers temporarily drop out of the economy, while the remaining ones eke some positive utility out of the brokers. Once a broker has a semi-coherent following, positive feedback sets in: the broker is encouraged by its clientele to provide articles that they will find interesting. Once the broker specializes to a small number of categories, other consumers are now attracted. By $t = 100,000$, all of the brokers have specialized into a single category, and no consumers receive any junk.

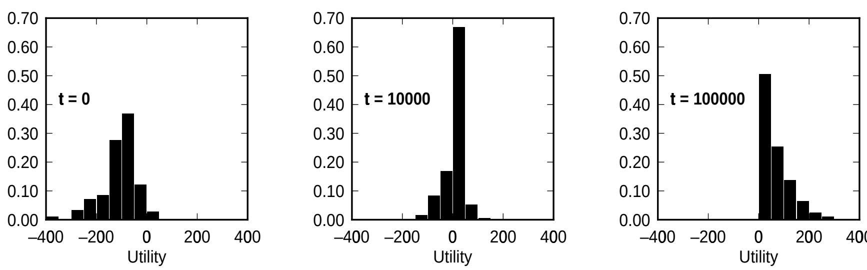

a) $t = 0$ b) $t = 10,000$ c) $t = 100,000$

Fig. 4. Evolution of the consumer utility distribution.

In this run, dramatic cyclical price wars were not observed, although more benign short-lived price wars played a role in driving superfluous brokers out of the market. Because the derivative-following price-setting algorithm forbids large discontinuous changes in prices, it appears unlikely to support dramatic cyclical price wars of the sort we found in the myoptimal case. Unfortunately, since the profit landscape typically contains several distinct humps separated by discontinuities, the derivative-following algorithm can often cause the economy to become stuck in highly suboptimal states. In other simulation runs in which the brokers initially offered an excessive number of categories, the system consistently evolved to a state in which no articles were sold. Making the system more robust to unfavorable initial conditions (or more responsive to environmental changes) requires permitting discontinuous jumps in price or interest vectors, but these capabilities also put the system at greater risk for cyclical price wars.

4 Conclusion

Our investigation of the dynamic behavior of an information filtering economy revealed at least two important effects: spontaneous specialization, which is generally desirable, and cyclical price wars, which are by and large undesirable even to consumers that may on the surface seem to benefit from lowered prices.

We found that specialization is driven by two distinct mechanisms working together. First, if the extrinsic transport and processing costs P_T and P_C are intermediate between the low-cost "spam" regime and the high-cost "dead" regime, a monopolist broker prefers to offer a small number of categories. Second, competition among multiple brokers encourages them to become monopolists in largely non-overlapping sets of one or a few categories. Niche specialization is typically desirable from the perspective of both the brokers and the consumers.

Standard models of price wars [18] typically lead to a *stable* point at which no one makes a profit. The news filtering economy is extremely prone to *unstable* limit-cycle price wars, behavior that can be traced to the multi-humped, discontinuous topography of the profit landscape. Price wars undermine the tendency of the system to efficiently self-organize itself. We found that cyclic price wars could be eliminated in a system of brokers and consumers that had little knowledge of the system state and very simplistic algorithms for updating prices and interests, permitting useful specialization to occur. However, one cannot conclude that individual ignorance leads to societal bliss. The conservative price-setting strategy makes the system less nimble, and more susceptible to failure. Furthermore, even if ignorance led to good collective behavior, it would hardly be a stable strategy: there would be a strong incentive to use a better informed or more intelligent agent that could outperform its weaker opponents. Other effects that may hinder price wars in human economies, such as explicit and tacit collusion, frictional effects, and spatial or informational differentiation are likely to be weaker in agent economies. Price wars may indeed prove to be a serious problem to contend with in large agent economies of any sort, and merit our continued attempts to understand and control them.

Acknowledgments

We thank Bill Schneider, Wenke Lee, Anastasia Anastasiadi, and Hoi Chan for their contributions to the information economy simulator and prototype.

References

1. W. P. Birmingham, E. H. Durfee, T. Mullen, and M. P. Wellman. The distributed agent architecture of the University of Michigan Digital Library. In *AAAI Spring Symposium on Information Gathering in Heterogeneous, Distributed Environments*. AAAI Press, 1996.
2. A. Chavez and P. Maes. Kasbah: an agent marketplace for buying and selling goods. In *Proceedings of 1st International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, London, April 1996.

3. S. Clearwater, editor. *Market based control: a paradigm for distributed resource allocation.* World Scientific, Singapore, 1995.
4. J. Eriksson, N. Finne, and S. Janson. Information and interaction in MarketSpace — towards an open agent-based market infrastructure. In *Proceedings of the Second USENIX Workshop on Electronic Commerce*, November 1996.
5. D. F. Ferguson. *The Application of Microeconomics to the design of resource allocation and control algorithms in Distributed Systems.* PhD thesis, Columbia University, 1989.
6. R. H. Guttman, A. Moukas, and P. Maes. Agent-mediated electronic commerce: A survey. *Knowledge Engineering Review*, June 1998.
7. J. E. Hanson and J. O. Kephart. Spontaneous specialization in a free-market economy of agents. In *Proceedings of Workshop on Artificial Societies and Computational Markets at the Second International Conference on Autonomous Agents*, May 1998.
8. T. Hogg and B. A. Huberman. Controlling chaos in distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 21:1325, 1991.
9. J. Hu and M. P. Wellman. Learning about other agents in a dynamic multi-agent system. In *Proceedings of the Second International Conference on Autonomous Agents*, 1998.
10. B. A. Huberman, editor. *The Ecology of Computation*, Amsterdam, 1988. North-Holland.
11. B. A. Huberman, R. M. Lokose, and T. Hogg. An economics approach to hard computational problems. *Science*, 1996.
12. J. O. Kephart, J. E. Hanson, and J. Sairamesh. Price-war dynamics in a free-market economy of software agents. In *Proceedings of Alife VI*, June 1998.
13. J. O. Kephart, T. Hogg, and B. A. Huberman. Dynamics of computational ecosystems. *Physical Review A*, 40:404–421, 1989.
14. J. O. Kephart, T. Hogg, and B. A. Huberman. Collective behavior of predictive agents. *Physica D*, 42:48–65, 1990.
15. J. F. Kurose, M. Schwartz, and Y. Yemini. A microeconomic approach to optimization of channel access policies in multiaccess networks. In *Proc. of 5th Int. Conf. Distrib. Comput. Syst.*, 1985.
16. M. S. Miller and K. E. Drexler. Markets and computation: agoric open systems. In B. A. Huberman, editor, *The Ecology of Computation.* North-Holland, Amsterdam, 1988.
17. M. Stonebraker et al. An economic paradigm for query processing and data migration in Mariposa. In *Proc. of Parallel and Distributed Information Systems*, pages 58–67, 1994.
18. J. Tirole. *The Theory of Industrial Organization.* The MIT Press, Cambridge, MA, 1988.
19. M. Tsvetovatyy, M. Gini, B. Mobasher, and Z. Wieckowski. MAGMA: an agent-based virtual market for electronic commerce. *Applied Artificial Intelligence*, 1997.
20. J. M. Vidal and E. H. Durfee. The impact of nested agent models in an information economy. In *Proceedings of the Second International Conference on Multi-agent Systems (ICMAS96).* MIT Press, December 1996.
21. M. P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1, 1993.
22. J. White. Mobile agents white paper. General Magic, Inc., 1996.

This article was processed using the LaTeX macro package with LLNCS style