Designs for a future microtask market

Dissertation proposal

M. Six Silberman

May 9, 2014

1 Goals

This dissertation has three goals. The first is to present actionable designs for a future microtask market. These designs should yield outcomes that workers and employers consider better than typical outcomes in Amazon Mechanical Turk (AMT), the paradigmatic microtask market. The second is to show how concepts often seen as vague or unoperationalizable in computing practice can be operationalized by drawing on the social science literature. In this case, I focus on fairness, power, and governance. The third is to show how putting these concepts into practice can improve outcomes in computing systems.

2 Motivation

This effort is motivated by the current state and apparent trajectory of crowd work research and practice.

'Crowd work' is a term that points to at least three things. These things are related but not the same. First, crowd work is an industry. That is, it is a specific set of people, relationships, organizations, artifacts, and practices. Second, crowd work is a way to get things done. That is, it is a set of ideas or techniques. Third, crowd work is an active research area in computing.

The crowd work industry is made up of companies such as Amazon, oDesk, Microtask, CrowdFlower, and 99designs, the people who work for them, the technologies they build and run, and the discourse they share through blogs, formal publications, and events such as CrowdConf.

Crowd work as a way to get things done is made up of a set of ideas and techniques. We can sketch this way of doing things with its language—the nouns and verbs that populate the discourse of its practitioners. The central verb in the world of crowd work is 'crowdsource.' Crowd employers 'crowdsource' work 'to' workers 'over' an online 'platform.'

There are different ways to crowdsource work. The most well-known are 'contest-based crowdsourcing' and 'microtasking.' Contest-based crowdsourcing involves holding a contest around a task. Many workers submit entries to the contest, but only one or

a few winners are paid. This model is popular in graphic design; the most well-known contest-based crowdsourcing platform is 99designs, a market for graphic design work.

In contrast to contest-based crowdsourcing, microtasking involves cutting work into small pieces and giving each piece to a worker. Two pieces of related or adjacent work might be given to the same worker, or to two workers on opposite sides of the world. These workers might never talk to each other. In fact, they might not know of each other's existence, except in the sense that each knows they are part of a large 'crowd' working on a job. Take for example the task of entering handwritten data from a scanned form into a database. A crowd employer might turn each field into a separate image, then turn the transcription of each field into a separate task. A worker in India might transcribe the last name field; a worker in Kenya the first name; and a worker in the United States the date of birth. Three microtasking platforms are Amazon Mechanical Turk, or AMT; Microtask; and MobileWorks. Some large technology companies run their own 'internal' microtask platforms. Microsoft's, for example, is called Universal Human Relevance System, or UHRS. There are different kinds of microtasking platforms. For example, AMT is a market, while MobileWorks assigns tasks algorithmically to workers. In this dissertation I propose to focus on microtasking, and specifically on microtask markets. As AMT is the paradigmatic microtask market, my model will be based on AMT.

The central generic nouns in the world of microtask markets are *requester*, *worker*, *platform*, *task*, *price*, and *time.* The central verbs are *post*, *choose*, *do*, *accept*, and *reject*. Requesters post tasks to a platform. Workers choose tasks. Foremost among the factors weighed by workers in choosing tasks is a task's price and how long it will take. After choosing a task, the worker does the task. After a worker does a task, the requester can accept or reject it. If the task is rejected, the worker is not paid—although the requester can still use the work submitted.

Crowd work is an area of active research in at least three fields of computing research: human computation, human-computer interaction, and computer supported cooperative work. Human computation is an approach to solving computational problems by incorporating human input into an algorithmic process. The field draws on a long tradition of organizing humans to do computation: before computers were silicon, they were human (e.g., Light 1999; Grier 2005). But in its current form, human computation grew out of artificial intelligence research in the middle of the first decade of the 2000s. On one hand, the most rigorous interpretations of "human computation" (e.g., Law and von Ahn 2011, pp. 3-5) do not include all examples of crowdsourcing. On the other, some interpretations are much broader, including all of crowdsourcing and a wide range of other phenomena (e.g., Michelucci, ed., 2013). But human computation as a subfield of computer science research and body of knowledge contributes significantly to crowd work and crowdsourcing practice. The full name, for example, of the main human computation conference—"HCOMP"—is "Conference on Human Computation & Crowdsourcing." If human computation as a subfield of computer science is "about" one thing in particular, it is crowdsourcing.

Human-computer interaction, or HCI, is an interdisciplinary field that emerged in the 1970s and 1980s. Its early focus was on practical problems such as designing aircraft cockpits to reduce (sometimes deadly) pilot error. Early HCI theory drew on computer science and psychology. The problems it addressed usually concerned one person

interacting with one computer. In the 1980s and 1990s, the focus of the field shifted to expanding access to computing, mainly by increasing its usability. One result of this shift was the development of the graphical user interface. Later developments include mobile computing (i.e., smartphones), interactive displays, and tabletop computing. By the mid 2000s, the effort to expand access to computing by increasing its usability had largely succeeded. By this time, a more broadly scoped strain of HCI research had developed. Researchers in this "third wave" of HCI base their work in frameworks from the social and cognitive sciences such as activity theory, phenomenology, and distributed cognition (e.g., Harrison, Tatar, and Sengers, 2007; Kaptelinin and Nardi 2012, pp. 1-2). And by the late 2000s, HCI researchers had begun to examine some of the unexpected, and perhaps unfortunate, consequences of the field's success. These researchers began to discuss new directions for the field in the context of the global information society the field had itself helped bring about. These discussions have led to a proliferation of new concepts, theories, and methods in the field.

Computer supported cooperative work, or CSCW, grew out of HCI in the 1980s. It was started by HCI researchers and computing practitioners making software for groups and organizations. CSCW has thus had a longer and more rigorous alliance with theory and method from the social sciences than HCI. But as HCI has broadened its scope, some researchers and research funders have seen fit to include both fields in a new category, "human-centered computing," or HCC.

People, papers, ideas, money, and specific working technologies travel easily between HCI and CSCW, and sometimes between HCI, CSCW, and human computation. Researchers in the three fields can often be found working in the same lab, department, or school. Yet they have their own emphases, bodies of knowledge, and, to some extent, values and cultures. They overlap and influence one another, but are distinct. As the recent re-categorization suggests, the overlap between HCI and CSCW is greater than the overlap between "human-centered computing" (i.e., HCI and CSCW together) and human computation.

Just as human computation, HCI, and CSCW are related and overlapping but not identical, the crowd work industry, crowd work as a set of ideas, and crowd work research are related and overlapping but not identical. People, specific working technologies, ideas, money, and practices do circulate between these realms. But industry and research have, at least ostensibly, their own rules. And "crowd work" as a model has been applied beyond both—for example, in the nonprofit sector.

Whether you think crowd work matters and is worth thinking and talking seriously about may depend on how popular you think it will get. One reasonable view was expressed by `spamgirl`, well-known crowd worker and organizer `spamgirl` in Amazon's Mechanical Turk system (AMT), in 2010. An interviewer asked her: "What do you think about the role of crowdsourcing in the future of employment?" She said:

> I see it as fairly insignificant. I don't think most work can be done this way, and I think the work that can be completed is, for the most part, too cheap for most to bother wanting to do it. (`spamgirl` 2010)

Considering crowd work from "inside" AMT in 2010, this view is easy to understand. It was, and still is, easy for requesters to make task design mistakes. It was, and still

is, hard to "learn the ropes" as a worker. It was, and still is, hard for workers and requesters to communicate about work. It was, and still is, hard for workers to find well-paying work. And it was, and still is, hard for requesters to find skilled workers and figure out how much to pay them.

But some of these problems are specific to current implementations (e.g., AMT) or endemic to the current way of thinking about microtask markets, which AMT has had a large role in shaping. None of these problems need be insurmountable. And, crucially, to the extent that they are general problems, many technologists in research and industry are hard at work solving them. The potential financial rewards are large. Organizations of all kinds see crowdsourcing as a way to cut costs on existing processes and cheaply expand operational capabilities. Researchers, especially in the social sciences, see it as a cheap way to get data. And the flexibility of crowd work appeals to many workers (e.g., Martin, Hanrahan, O'Neill, and Gupta 2014). Technologists and business people call crowdsourcing the "future of labor." New crowdsourcing startups appear with high frequency. And the volume and scope of crowd work research is growing. Major HCI and CSCW conferences devote multiple sessions to the topic, and it has a new conference of its own. From within this context, Kittur, Nickerson, Bernstein, Gerber, Shaw, Zimmerman, Lease, and Horton, in a paper published at the CSCW conference called, simply, "The future of crowd work" (2013), write:

> While not all jobs are amenable to being sent down a wire, there are portions of almost any job that can be performed by the crowd. We foresee a world in which crowd work continues to expand...

More recently still, Cefkin, Anya, Dill, Moore, Stucky, and Omokaro, in a summary of a workshop held at the CSCW conference called "Back to the future of organizational work: crowdsourcing and digital work marketplaces" (2014), write:

> Businesses increasingly accomplish work through innovative sourcing models that leverage the crowd... It is our position that crowdwork is likely to be increasingly integrated into existing ways of doing organizational work.

All that is needed for crowd work to keep growing is for it to keep making, or saving, employers money. It is doing so now. Given the practical and intellectual effort going into crowd work, it seems likely to do so for some time. But this growth says little about workers' experience. The lived experiences reported by crowd workers do not match up to the breathless excitement of technologists, managers, and pundits hailing crowd work as the future of labor. Organizational incentives are such that technologists are often allocated to make sure workers can work effectively. But often, "effectively" means "just effectively enough."

Crowd work matters because it is likely to grow. As it grows, it will play a bigger role in the lives of more and more people. It will do this largely by playing a bigger role in work arrangements in many fields. This growth will create both new opportunities and new expectations. It will "disrupt" existing business strategies and livelihood strategies for both organizations—potential crowd employers—and employees—potential crowd workers. The number of people who make a living or second job through crowd work will grow.

There are at least two big open questions about this change. The first is the distribution of benefits. If current trends hold, the distribution of benefits will be starkly uneven. At present, crowd employers benefit tremendously from greatly reduced costs and expanded operational capabilities compared to traditional employment arrangements. But these benefits are secured, to a large extent, at workers' expense. This is not to say that workers do not value the unique opportunity crowd work affords. They do. Specifically, they value the flexibility it affords—the opportunity to work from home, on their own hours. But this flexibility has a cost (e.g., Silberman, Ross, Irani, and Tomlinson 2010; Irani and Silberman 2013). Crowd work often pays less than similar work within a traditional employment relation. Legally, crowd workers are "independent contractors." As such, they are not entitled to minimum wage, group health insurance, or overtime pay. And crowd work pay is uncertain. In a microtask market such as AMT, a requester can reject work at any time, for any or no reason. Workers are not paid for rejected work. And they have few clues about how likely this is to happen for a given task. The ability of requesters to refuse to pay for submitted work—work they may keep and use—with impunity has adverse consequences for both workers and requesters. For example, it makes workers' livelihoods vulnerable to predictable technical or administrative errors without encouraging employers to take reasonable steps to prevent them. This risk incentivizes workers to spend less time on any given task. Less time on task often means inferior work. Large amounts of inferior work means employers must develop sophisticated quality control strategies. This raises the effective cost of crowdsourcing. But this cost does not benefit anyone: it is just lost time. And it contributes to a common view among employers and researchers of crowd workers as low-skill, interchangeable, and untrustworthy. Aggressive quality control strategies lead workers to see many employers as miserly, nitpicking, or even cheating. Even from an entirely financial perspective, this adversarial environment can be shown to be inferior to more cooperative arrangements that are easily imagined. Indeed this and many other adverse consequences of the current approach to microtask market design are practical, quantifiable, explicable, and even to some extent predictable.

Finally, workers are often invisible to employers, in two senses. First, employers know very little about the personal constraints, working conditions, or needs of workers. AMT especially encourages this anonymity by representing workers as alphanumeric strings. Second, and perhaps more fundamentally, many crowd employers do not think of workers as people at all. After presenting "Being a Turker" at the CSCW conference in 2014, David Martin (pers. comm., 2014) wrote:

> Our presentation seemed to go pretty well at CSCW. Some people even came up to me after to say, "I had never really thought before that using MTurk involved real people, and that you maybe should pay them a decent amount..."

In my experience, this is a relatively common response from crowd employers to questions about workers' rights. AMT in particular fosters this ignorance through its language and interface design (Silberman, Irani, and Ross 2010; Irani and Silberman 2013), but it is not unique to AMT.

The second big open question about crowd work is harder to address quantitatively. The philosopher Donna Haraway (Nakamura 2003) asks it pithily: "What kind of world

is this?" Put as a question to computing researchers, crowd employers, and crowd work platform operators, the question is: What kind of world are we building? To ask this question is not to argue that workers do not value crowd work. They do. But if we pay close attention to their discourse (e.g., `spamgirl` 2010; Martin, Hanrahan, O'Neill, and Gupta 2014), we see that they participate in it much in the same way that I would agree to give someone $50 for use of a ladder after they had pushed me into a deep hole. That is, workers seem to participate in crowd work not with the excitement of participating in a gleaming new technological future, but wearily, with resigned acceptance. There is no straightforward or universally unobjectionable way to talk about this. But it seems irresponsible to ignore it.[1]

At stake, ultimately, is the future of crowd work—and perhaps, if the pundits are to believed, the future of work. Like Kittur et al. (2013), I expect crowd work to continue to grow. Employers and technologists will respond to the great institutional, economic, and cultural incentives for technological and organizational innovations that drive down labor costs and expand operational capabilities. Workers will respond to the opportunity to work in a time and place of their choice. In this context, crowd work arrangements that support viable livelihoods could indeed unlock "an incredible number of opportunities for careers in skilled work" (Kittur et al. 2013, p. 1301). Such growth could contribute to socioeconomic recovery in the recession-hit North, sustainable socioeconomic development in the South, and socioeconomically beneficial globalization. Human computation and HCC research has already begun to explore the potential benefits—for employers and workers—of bringing crowd work to "developing regions," "the bottom of the pyramid," and "low-income workers" (e.g., Khanna et al. 2010; Narula et al. 2011; Gawade et al. 2012; Gupta et al. 2012; Samdaria et al. 2012). Yet Kittur et al.'s pessimistic scenario, in which crowd work falls "into an intellectual framing focused on low-cost results and exploitative labor," with workers assumed "interchangeable and untrustworthy, [with] low or static skill sets and strong motivations to shirk" (2013, p. 1301), is easier to imagine—in part because it is already realized in AMT, the dominant and paradigmatic microtask market.

We are on a bad road in crowd work. This dissertation is about how to get off it.

3 Method

I propose to achieve the goals of the dissertation using agent-based modeling as an illustrative and experimental method.

Agent-based modeling is widely used in ecology (e.g., Janssen 2002; Grimm and Railsback 2005; Janssen and Ostrom 2006), sociology (e.g., Macy and Willer 2002), economics (e.g., Schelling 1978; Testfatsion and Judd 2006), and interdisciplinary social science (e.g., Bonabeau 2002; Epstein 2006). Agent-based models represent phenomena as "evolving system[s] of autonomous interacting agents" (Janssen and Ostrom 2006). The method is suited "for studying systems [...] composed of interacting agents

[1]The analogy above is owed to the economist Randall Bartlett, specifically his 1989 book *Economics and Power*, on which I will draw heavily in the dissertation. That such an illuminating analogy should surface in a book on economics and power is, on reflection, unsurprising: a robust discourse on the role of power in economic life is exactly what computing culture lacks.

[that] exhibit emergent properties" (Tesfatsion and Judd 2006). Such properties result from agents' interactions and are not predictable with knowledge of the agents alone.

Crowd work markets, like the markets studied by the contributors to the *Handbook of Computational Economics Vol. 2: Agent-based Computational Economics* (Tesfatsion and Judd 2006), are just such complex systems. Workers, employers, and market operators all have some behaviors and needs in common but are diverse among many axes. And no single mechanism explains the large-scale patterns—e.g., the relative distribution of "gains from trade" between workers and employers—associated with particular market designs. Rather, such patterns "emerge" from the repeated interactions of the agents that make up the market. Agent-based modeling is thus an appropriate method for exploring the dynamics of current crowd work markets and the design space of possible future markets.

I will model a *market*[2] with two kinds of agents, *workers* and *employers*, and one kind of object, *tasks*. Tasks are generated exogenously and first appear in an employer's *task queue*. Tasks appearing in an employer's task queue have a *value*. This is the value to the employer of having the task *completed*. I assume employers do not complete tasks themselves but instead *post* them to the *market queue* for workers to complete. Before posting a task to the market queue, the employer must give the task a *price*. This is the price the employer will give to the worker who completes the task. The price will usually be less than the value, but some employers may price tasks high with the a priori intent to refuse payment. Workers select tasks in a random order. After a task is completed, the employer decides whether to *pay* the worker the listed price for completing it or to *refuse* payment.

I will first use this model to explore the basic dynamics of supply and demand in the market. I will then explore how design variables such as the option to reject work; reputation; and agent communication affect outcomes such as posted task prices; the distribution of the gains from trade; and worker perceptions of crowd work as a viable livelihood. I list design and outcome variables below.

LIST 1. DESIGN VARIABLES.

- the option to "reject" (i.e., not pay for) work
- use of a "majority rule" scheme for automating rejection decisions
- a clear set of criteria for rejecting work
- a mediated rejection "appeal" process
- a worker reputation system
  - based on approval (i.e., payment) rate
  - based on references from employers
  - based on data about past work
- an employer reputation system
  - based on references from workers
  - based on data about past tasks (e.g., posted prices, nonpayment history, posted vs. actual completion times)
- integration of employer reputation into the task browse/search interface
- worker and employer profiles

[2] In this paragraph and the next, nouns and verbs in italics will be represented explicitly in the simulation.

- anonymity
- support for worker collaboration
- support for worker-employer communication about tasks

LIST 2. OUTCOME VARIABLES.

- posted task prices
- posted task time estimates
- employer-estimated wage
- workers' a priori wage estimates
- workers' a priori estimates of payment probabilities
- actual payment probabilities
- actual wages
- ratios of wages to costs of living
- total value of tasks posted
- distribution of the gains from trade between stakeholders
- employers' estimates of worker probability of shirking
- actual probability of shirking
- prevalence of deceptive employer behavior (e.g., posting a task with a high price with the intention to reject a large fraction of submitted work)
- perceptions of trustworthiness (i.e., estimates of others' strategies)
- workers' perceptions of crowd work as a viable livelihood
- work quality
- the distribution of skills among workers

After exploring the dynamics of the market and the relation between design and outcome variables, I will draw on literature in the social sciences to introduce notions of fairness (e.g., Fehr and Schmidt 1999), power (e.g., Bartlett 1989), and governance (e.g., Ostrom 1990, 2005) to the model. These chapters will show how these concepts can be operationalized to guide both the evaluation of market outcomes and the design of the market. Finally, I will consider possible implications of operating a microtask market through a distributed protocol rather than as a monolithic platform under administrative control of a single organization.

4 Audience

This dissertation's main intended audience is researchers in human computation and human-centered computing (HCC). The ideas presented in it may also be of interest to crowd work practitioners (workers, employers, and platform operators), although the ideas may need to be presented in a different format and venue to come to practitioners' attention. The dissertation may also be of interest to researchers in science and technology studies and the social sciences.

5 Contribution

This dissertation will offer two main contributions to human computation and HCC research. The first is the series of implementable designs for new microtask markets. These designs will yield outcomes users will recognize as "better" than those achieved by AMT. This will be achieved mainly by incorporating considerations of "human factors"—fairness, power, and governance—thus far underexplored in crowd work research in human computation and HCC. The second is the operationalization of these concepts through reference to the social science literature. As computing systems have grown ubiquitous, they have come to mediate more and more human interaction, including the livelihoods of many information workers. With this growth, the ideas about social and economic life that shape these systems and their day-to-day operation play an ever-larger, if often invisible, role in shaping the lives of more and more people. In crowd work, these ideas come mainly from two discourses: computing and business. To the extent that these discourses focus on the "rational"—i.e., short-sightedly self-interested—aspects of social and economic life, crowd work systems will continue to encourage short-sightedly self-interested behavior from participants, with suboptimal results for many over the long term. To the extent, on the other hand, that concepts from contemporary social science that offer a richer view of social and economic life can be brought into dialogue with the practical concerns of software design, operation, and maintenance, such dialogue will expand, for the better, the kinds of software we can imagine and build, and the ways we can interact with one another through it.

6 Structure

The dissertation will have three parts. Part I will have three chapters. Chapter 1 will introduce crowd work and microtask markets. Chapter 2 will introduce agent-based modeling. Chapter 3 will develop the "basic" agent-based microtask market model. Part II will include four chapters, each covering one concept. Chapter 4 will discuss fairness; Chapter 5, power; Chapter 6, governance; and Chapter 7, locus of control. Chapters 4–6 will each have four sections. The first section of each chapter will introduce the concept based on a selection of relevant social science research. The second will add the concept to the basic market model. The third will explore changes to the basic design suggested by the concept. The fourth will assess how these changes affect market outcomes. Chapter 7, on locus of control, will focus on one question: what are possible implications of developing a microtask market *protocol* rather than a new microtask market under the control of a single organization? Part III will include three final chapters. Chapter 8 will synthesize the preceding analyses. Chapter 9 will present several designs for a future microtask market. Chapter 10 will conclude the dissertation with a discussion of computing research and practice, social value, and social science.

References

Bartlett, R. 1989. *Economics and Power*. Cambridge.

Bonabeau, E. 2002. Agent-based modeling. *Proc. Natl. Acad. Sci.* **99**(S3): 7280-7287.

Cefkin, M., O. Anya, S. Dill, R. Moore, S. Stucky, O. Omokaro. 2014. Back to the future of organizational work. *CSCW '14 Companion*: 313-316.

Epstein, J. M. 2007. *Generative Social Science*. Princeton.

Fehr, E., K. M. Schmidt. 1999. A theory of fairness, competition, and cooperation. *Quarterly J. Econ.* **114**(3): 817-868.

Gawade, M., R. Vaish, M. Nduta, J. Davis. 2012. Exploring microwork opportunities through cybercafés. *DEV '12*: 21.

Grier, D. A. 2007. *When Computers Were Human.* Princeton.

Grimm, V., S. F. Railsback. 2005. *Individual-based Modeling and Ecology*. Princeton.

Gupta, A., W. Thies, E. Cutrell, R. Balakrishnan. 2012. mClerk. *CHI '12*: 1843-1852.

Harrison, S., D. Tatar, P. Sengers. 2007. The three paradigms of HCI. *alt.chi '07*.

Irani, L., M. S. Silberman. 2013. Turkopticon. *CHI '13*: 611-620.

Janssen, M. A., ed. 2002. *Complexity and Ecosystem Management*. Edward Elgar.

Janssen, M. A., E. Ostrom, eds. 2006. Special issue: Empirically-based, agent-based modeling of social-ecological systems. *Ecology and Society* **25**.

Kaptelinin, V., B. Nardi. 2012. *Activity Theory in HCI*. Morgan & Claypool.

Khanna, S., A. Ratan, J. Davis, W. Thies. 2010. Evaluating and improving the usability of Mechanical Turk for low-income workers in India. *DEV '10*: 12.

Kittur, A., J. V. Nickerson, M. S. Bernstein, E. M. Gerber, A. Shaw, J. Zimmerman, M. Lease, J. J. Horton. 2013. The future of crowd work. *CSCW '13*: 1301-1318.

Law, E., L. von Ahn. 2011. *Human Computation*. Morgan & Claypool.

Light, J. S. 1999. When computers were women. *Technology and Culture* **40**(3): 455-483.

Macy, M. W., R. Willer. 2002. From factors to actors: computational sociology and agent-based modeling. *Ann. Rev. Soc.* **28**: 143.

Martin, D., B. V. Hanrahan, J. O'Neill, N. Gupta. 2014. Being a Turker. *CSCW '14*: 224-235.

Michelucci, P., ed. 2013. *Handbook of Human Computation*. Springer.

Nakamura, L. 2003. Prospects for a materialist informatics. *Electronic Book Review*, 30 Aug 2003.

Narula, P., P. Gutheim, D. Rolnitzky, A. Kulkarni, B. Hartmann. 2011. MobileWorks. *HCOMP '11*.

Ostrom, E. 1990. *Governing the Commons*. Cambridge.

Ostrom, E. 2005. *Understanding Institutional Diversity*. Princeton.

Samdaria, N., A. Mathur, R. Balakrishnan. 2012. Paying in kind for crowdsourced work in developing regions. *Pervasive '12*: 343.

Schelling, T. 1978. *Micromotives and Macrobehavior*. Norton.

Silberman, M. S., L. Irani, J. Ross. 2010. Ethics and tactics of professional crowd work. *XRDS* **17**(2): 39-43.

Silberman, M. S., J. Ross, L. Irani, B. Tomlinson. 2010. Sellers' problems in human computation markets. *HCOMP '10*: 18-21.

`spamgirl`. 2010. Spamgirl speaks! An interview (part 3 of 3). *Broken Turk*, Nov 2010.

Testfatsion, L, K. L. Judd, eds. 2006. *Handbook of Computation Economics Vol. 2: Agent-Based Computational Economics*. North-Holland/Elsevier.