Electronic Commerce Research, 1: 239–263 (2001)

Designing Agent-Based Electronic Employment Markets

WILLIAM R. GATES and MARK E. NISSEN {brgates; MNissen}@nps.navy.mil
Graduate School of Business and Public Policy, Naval Postgraduate School, Monterey, CA 93943, USA

Abstract

Two modes of matching people with jobs prevail at present: 1) hierarchical planning and 2) distributed markets. Each has strengths and limitations, but few systems have been designed to take advantage of strengths corresponding to both. With evolving information technology, however, the job-matching process could be accomplished far more equitably and efficiently using web-based markets within the firm, and intelligent agents offer excellent potential to help both potential employees and employers find one another in a distributed, electronic marketplace. But realizing this potential goes well beyond simply changing the rules of internal job matching or making agent technology available to job searchers. Rather, the corresponding markets and technologies must be *designed*, together, to mutually accomplish the desired results (e.g., efficient and effective matching) and conform to necessary properties (e.g., market clearing). Through the research described in this paper, we draw from Game Theory results to assess the feasibility of using two-sided matching algorithms to address this market-design problem. We also draw from current agent research to address the information technology dimension of the problem by implementing a proof-of-concept multi-agent system to enact, automate and support the corresponding market solution. This paper integrates the key economic and technological elements required to design robust electronic employment markets. And the corresponding research provides new knowledge and insight into co-development of the requisite economic markets and agent technologies.

Keywords: intelligent agents, multi-agent systems, web-based markets, electronic commerce, two-sided matching markets, game theory, electronic labor markets

1. Problems with current employment approaches

Two modes of matching people with jobs prevail at present: 1) hierarchical planning and 2) distributed markets. Patterned after centrally-planned (e.g., former Soviet-style) economies and command-and-control (e.g., military) organizations, the former approach remains prevalent for matching job candidates to jobs *within* the current enterprise. As an example from the U.S. military, the Navy currently matches sailors to jobs using a centralized, labor-intensive detailing process, one which leaves many parties (e.g., sailors, commands) dissatisfied and results in poor employee morale, performance and retention.

Alternatively, the latter, market-based approach supports unrestricted, point-to-point matching between potential employees and *outside* employers. As an example from the information technology (IT) environment, technology professionals in nearby Silicon Valley currently have access to a hyperactive job market—characterized by negative unemployment—in which a multitude of job opportunities is available to many people. In this situation, information overload—for example, associated with the requirement to search through, screen and filter myriad job opportunities—has become problematic, and employee turnover is now incessant.

With evolving information technology, the job-matching process could be accomplished far more equitably and efficiently using web-based markets within the firm, and intelligent agents offer excellent potential to help both potential employees and employers find one another in a distributed, electronic marketplace. But realizing this potential goes well beyond simply changing the rules of internal job matching or making agent technology available to job searchers. Rather, the corresponding markets and technologies must be *designed*, together, to mutually accomplish the desired results (e.g., efficient and effective matching) and conform to necessary properties (e.g., market clearing).

Through the research described in this paper, we draw from Game Theory results to assess the feasibility of using two-sided matching algorithms to address this market-design problem. And we draw from current agent research to address the information technology dimension of the problem by implementing a proof-of-concept multi-agent system—called the "Personnel Mall"—to enact, automate and support the corresponding market solution. The balance of the paper follows this introduction by discussing employment market economics, after which we provide an overview of intelligent agent technology and describe the Personnel Mall in some detail. We then integrate the key economic and technological elements required to design robust electronic employment markets. The corresponding research provides new knowledge and insight into co-development of the requisite electronic markets and agent technologies.

2. Labor market economics

As background, we draw from what is now textbook understanding of labor market economics (cf. Ehrenberg and Smith, 1997). This discussion begins with an overview of market-based labor markets, which provides the basis for comparison and contrast with hierarchical labor markets. Then moving into current economics research, we discuss two-sided matching games to provide an understanding of the mechanisms under consideration for incorporation into agent-based labor markets.

2.1. Market-based labor markets

Market-based approaches to employee/employer matching rely on the interaction of labor demand and supply (Ehrenberg and Smith, 1997). Figure 1 illustrates labor demand and supply curves for a representative labor market (e.g., manufacturing, software development, clerical support). The wage rate in this market tends toward its equilibrium value (W^*), where the demand and supply curves intersect. The quantity of labor that employers willingly hire at this wage rate exactly equals the quantity of labor that employees willingly supply (L^*)[1]. Anyone that wants to work in the industry can find sufficient work, and any firm that wants to hire employees can find adequate employees.

If the wage rate is above its equilibrium value, there is an excess supply of labor. At the higher wage rate, employers choose to hire less labor (e.g., by reducing output, re-

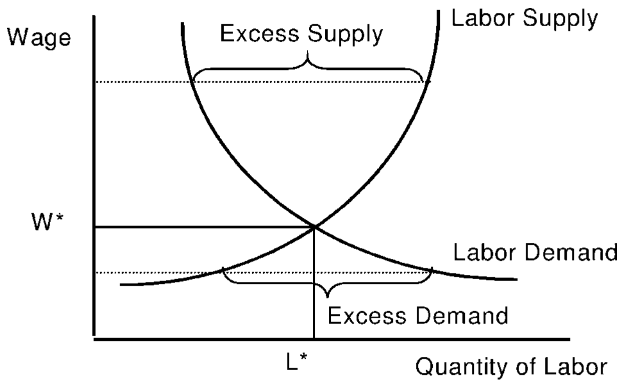

Figure 1. Market-based labor markets.

placing labor with capital), while employees choose to supply more labor (e.g., through longer workweeks, new entrants to the industry's workforce). Through these dynamics, competition for jobs will force the wage rate to fall until the quantity of labor demanded equals the quantity of labor supplied (i.e., the equilibrium wage rate). If the wage rate is below equilibrium, opposite forces will create an excess demand for labor; competition for employees will increase the wage rate to its equilibrium value.

A subtle but important aspect of equilibrium wages rates involves job amenities (e.g., work environment, geographic location, commute, promotion potential, work content/challenge, job satisfaction). In weighing employment benefits in one industry relative to alternative time uses (e.g., leisure and other jobs), job amenities are important considerations. If job amenities are particularly attractive in one industry, individuals will supply labor to that industry at relatively low wage rates; if job amenities are unpleasant, labor is only supplied at relatively high wage rates (Ehrenberg and Smith, 1997; Moore and Viscusi, 1990). Compensating wage differentials are illustrated by comparing wage rates either across industries (e.g., wage rates for schoolteachers versus garbage collectors or coal miners) or for a particular industry across geographic locations or work environments. Holding other characteristics constant, individuals willing to work in an industry for relatively low wages either receive particularly high benefits from agreeable job amenities or are relatively weakly deterred by objectionable job amenities.

Further developing labor demand and supply provides additional insight into market-based labor markets. Labor demand is determined by the value employers derive from hiring an additional employee (holding the levels of all other inputs constant). In for-profit businesses, the market value of the extra output produced measures labor's contribution. It pays to hire an additional employee as long as the value of the output the employee produces exceeds the cost of hiring that employee (i.e., the marginal value of labor exceeds the marginal cost of labor). Thus, firms hire employees if they value labor at or above the market wage rate; firms (uses) that generate values below the market wage rate choose not to hire labor. Market-based labor markets also allocate labor to its highest valued (priority) uses; as the market wage increases, lower valued uses dropout while higher valued uses remain filled.

The labor supply in a particular industry reflects the relationship between the labor employees are willing to supply and the market wage rate. The wage rate implied here reflects total compensation including, salary, benefits, bonuses and others. Employees' decisions regarding the labor they will supply for a particular industry encompass both job amenities and the benefits employees receive from the best alternative use of their time, including either their value of leisure or the benefits forgone by not working in other industries. As wages increase for the industry in question, the quantity of labor supplied typically increases. At higher wage rates, the benefits of working in this industry increase relative to the value of both leisure and employment in other industries. Individuals already working in the industry may choose to work longer hours (if institutionally possible), and new employees will choose to enter the industry (e.g., by entering the workforce, leaving other industries). Thus, market-based labor markets ensure that the employees hired are the most willing to work in the industry. Anyone can find work that is willing to work at or below the market wage rate; individuals demanding higher wages choose to use their time in other ways. As market wages decrease, individuals less willing to work voluntarily leave (e.g., because they have better alternative options, they are relatively less attracted (more put off) by the job's amenities); those most attracted to the industry remain.

Market-based labor markets balance demand and supply, ensuring equality between the quantity of labor demanded and supplied. However, market-based labor markets also allocate labor to its highest valued uses (i.e., demand efficiency) and to the uses for which it is best suited (i.e., supply efficiency). The demand and supply diagram depicted in Figure 1 only considers a single industry (use) for labor (i.e., partial equilibrium). This tends to emphasize the balance between demand and supply; but it does not specifically illustrate that labor has other uses. As the analysis expands to multiple industries/labor markets (i.e., general equilibrium), the importance of demand and supply efficiencies becomes more evident. However, the information requirements in market-based labor markets also become more evident. To operate efficiently, employees must have complete information about all relevant job opportunities, including salary, benefits and job amenities (e.g., work environment and content, promotion potential, commute). In the context of labor markets with national or global reach, clearly, some kind of technology is required to help people manage the corresponding abundance of labor-market information.

2.2. Hierarchical labor markets

Hierarchical labor markets assign individuals to jobs using a centralized process. Examples of hierarchical labor markets include job assignments within commercial firms, government agencies and the military's labor detailing process. Hierarchical job assignments must rely on administrative procedures to match individual capabilities and job requirements and to reflect both the job's relative priority and the individual's job preferences. There is no mechanism to automatically strike a balance between supply and demand efficiencies, as in market-based labor markets. At one extreme, employers can assign individuals to jobs with little regard to personal preferences. Employees can either accept the assignment or find an alternative occupation. This approach emphasizes the employer's performance (i.e., demand efficiency) at the expense of employee morale (i.e., supply efficiency). At the other extreme, employers can emphasize individual job preferences relative to job priority and the match between employee skills and job requirements. This has the reverse emphasis (i.e., supply over demand efficiency). Criticisms against hierarchical labor markets involve their slow reaction time to shifts in labor demand or supply and their inability to ensure demand and supply efficiencies. These represent inherent equilibrium conditions in market-based labor markets. Such criticisms reflect cumbersome administrative employee/job matching procedures, intensive information requirements and asymmetric incentives (e.g., profits vs. morale).

Demand and supply efficiencies are particularly important for closed systems with a constrained labor supply, such as commercial firms where labor requires significant firm-specific knowledge or the military services and other, specialized government agencies. When labor requires employer-specific skills, it is difficult to hire mid-career employees to fill areas of need; employers develop an internal labor force through education, training and promotion, and allocate this labor force across job vacancies. In the military and many large commercial firms, wages are relatively uniform across jobs requiring similar skills and experience (i.e., no compensating wages) and "sticky" (i.e., slow to change). As a result, the cost of assigning labor to one use is the loss of output in the best alternative unfilled use for that labor (i.e., opportunity cost); salaries and benefits are irrelevant in measuring labor costs. If labor assignments do not maximize demand and supply efficiencies, the system wastes resources by applying them to less valuable jobs. And it reduces job satisfaction, morale and retention by assigning labor to jobs that are relatively less desirable with no compensating wage differential.

2.2.1. Navy example To further illustrate the issues involved with hierarchical labor markets, it is useful to briefly describe the U.S. Navy's enlisted personnel distribution system. Because of the Navy's large size, global presence, unique mission and policy of frequent employee job rotation, it is extremely difficult to achieve efficient employee/job matching in this system. The extreme nature of the Navy system makes it particularly attractive for research; that is, if problems with such an extreme system can be solved, then the results of this study should also generalize quite well across many corporations, government agencies and other military branches.

The Department of the Navy (DoN) uses a centralized, hierarchical labor market to match enlisted sailors to jobs (U.S. Navy, 2000). On the demand side, Navy commands (e.g., ships, bases) identify open positions. Job vacancies are compared to projections of available personnel. Typically, the number of positions to be filled exceeds the supply of available personnel. Therefore, the Navy develops its Manning Plan that spreads the labor shortage across all commands, on a "fair-share" basis. The Navy then prioritizes job vacancies based on each command's mission, current staffing levels, and several other relevant characteristics. This process attempts to distinguish between high and low valued demands for labor, to mimic demand efficiency in market-based labor markets.

On the supply side, available personnel are categorized according to their qualifications (i.e., ratings), including skills, experience, education/training, career path and others. Similar skill groups are arranged in "communities" (e.g., electronics, supply, machinists). Each community has a detailer charged with matching personnel to jobs. Sailors seeking job assignments can express their personal preferences to the detailer. The detailer is responsive to job vacancy priority ratings, but there is some room for discretion in tailoring job assignments to meet the sailors' personal preferences (i.e., supply efficiency). Supply efficiency is, clearly, subordinate to demand efficiency in this process, however.

DoN's hierarchical labor market is further complicated, because enlisted sailors change jobs every two to three years. Thus, the centralized detailing process reassigns between one third and one half of the enlisted force (e.g., 100,000–150,000 people) every year. This adds a magnitude and time dimension to the process that is more critical than in typical civilian labor markets. The Navy begins identifying job vacancies and available personnel as early as nine months in advance of planned rotation. Time also affects the job vacancy priority rating. More imminent vacancies receive a higher priority than similar but more temporally-distant vacancies.

From this brief introduction, it is clear that DoN's centralized detailing process has developed administrative mechanisms to try balancing the quantity of labor supplied and demanded, as well as demand and supply efficiency. DoN fills billets (i.e., jobs) according to a predetermined priority ranking until the labor supply is exhausted, and demand efficiency is emphasized over supply efficiency. In contrast, with market-based labor markets, equilibrium wage rates automatically perform these functions; wages adjust until there is no excess supply or demand for labor, and employees voluntarily choose their preferred job, considering both relative wages (i.e., compensating wage rates) and job amenities. In DoN's hierarchical labor market, wage rates do not increase to limit the demand for labor to the available supply, so commanding officers are frustrated they cannot fill vacant positions. Similarly, wages do not adjust across job assignments to account for job amenities, and assignments do not fully incorporate the sailor's job preferences. Predictably, both commanding officers and enlisted sailors voice dissatisfaction with the current hierarchical labor market.

2.3. Two-sided matching markets

Unlike fast-paced IT firms in Silicon Valley, wage rates for military personnel—and most large corporations—are set by fiat (e.g., by Congress, the Personnel Department) and adjust very slowly to supply- or demand-driven pressures. At least in the short term, the Navy—and most other large organizations—cannot rely on spot labor markets for filling its key jobs with qualified people. Indeed, without its current, hierarchical detailing system, the Navy would find it very difficult to fill many of its important jobs. Yet the Navy—as well as other major enterprises—could also benefit from the efficiencies associated with market-based systems. This conundrum leads us to draw from Game Theory and consider two-sided matching markets (Roth and Sotomayor, 1990). Such matching markets offer some benefits of market-based markets while retaining administrative control through a centralized organization. But even such market hybrids involve huge information requirements which call for some kind of IT support (e.g., electronic labor markets).

A two-sided matching market assigns individuals to jobs when there are several possible employers and employees. The matching algorithm balances the employers' and employees' preferences, but it can produce assignments that give priority to either employers or employees. As such, the algorithm specifically addresses *both* demand and supply efficiency. Two-sided matching algorithms are currently used in assigning medical students to residency programs (Roth and Sotomayor, 1990; Roth, 1984; Roth and Peranson, 1997) and pledges to sororities at some colleges and universities (Mongell and Roth, 1991).

2.3.1. Medical residency example The market for medical residents illustrates the two-sided matching system. As U.S. students complete their final year of medical school, they interview for residency positions. Each student interviews with several residency programs, and each program interviews several students. After the interviews, students rank residency programs according to their individual preferences, and programs independently rank students according to their preferences. Students and programs submit their prioritized lists to a central clearinghouse. The clearinghouse compares the lists and assigns students to programs. On a predetermined date, students and residency programs receive their assignments. Each matched student is assigned to one residency program, and each program is assigned students up to the number of available positions. Unmatched students individually seek unfilled positions; programs with unfilled positions can seek either unmatched U.S. medical students or foreign-trained students.

Participating in this centralized assignment process is voluntary. Residency programs and medical students are free to establish individual agreements, but over 90% of assignments are made through this voluntarily, centralized process. To generate this participation level, the matching process must satisfy a few basic conditions. One of the most important conditions is *stability*: both students and programs must be at least as happy with their assigned match as with any agreement they could reach outside the centralized process. The outcome is unstable if a student and unmatched program both prefer one another to the respective program and student with which they are centrally matched. With unstable matches, the student and program would both choose to forgo the assigned match and form their own side agreement. If a student is not matched to his or her highest ranked program, the program must have been assigned students that it ranked more highly (i.e., the program would not reject the assigned match). Conversely, if a program does not receive its highest ranked students, these students must be matched with programs they rank more highly (i.e., the students would not reject the assigned match). For descriptions of the problems encountered when programs do not meet these requirements, see (Roth, 1991).

2.3.2. One-to-one matching example A number of alternative matching algorithms are available to implement two-sided matching markets. Describing one, simple matching algorithm helps illustrate the issues involved in designing electronic employment markets. Suppose each company requires one employee to fill one job vacancy, and each employee can only fill one job vacancy. Furthermore, suppose salaries are predetermined and invariant for each vacancy (e.g., as in the large corporation, government agency, military service). This one-to-one matching process is equivalent to the “marriage market,” which has been studied extensively (cf. Gale and Shapley, 1962; Crawford, 1998; Harrison and McCabe, 1989; Roth and Sotomayor, 1990, pp. 15–121). In this model, salary can be considered a job characteristic that affects employees’ job preferences, along with work environment, promotion prospects, job location/commute and others. This simplified model provides a convenient starting point for market design.

In the one-to-one matching model, an iterative matching algorithm begins when employers identify and rank all candidates qualified to fill a job vacancy. Each employer then extends a job offer to its highest-ranked candidate. Candidates with multiple offers tentatively accept the job proposal they most prefer; candidates reject offers they consider unacceptable (i.e., they would rather not work than accept that offer). Each employer that has a job vacancy after this first round extends an offer to its second highest-ranked candidate, whether or not that employee has entered a tentative agreement. Again, candidates tentatively accept their most-preferred offer (i.e., potentially rejecting offers accepted in the previous round) and reject unacceptable proposals. The process continues through additional rounds until no tentative agreement or newly tendered offer is rejected. After completing this process, some job vacancies may remain unfilled (e.g., if all acceptable employees are assigned to jobs the employees prefer to the unfilled vacancies). Similarly, some personnel may be unemployed (e.g., if all acceptable jobs are filled with candidates preferred by the employers). This provides a contrast to the market-clearing properties of market-based labor markets discussed above.

Nonetheless, Roth and Sotomayor (1990) have shown that this matching algorithm produces a stable outcome. However, this algorithm produces an outcome that emphasizes employer relative to employee preferences (i.e., an “employer-optimal” outcome that emphasizes demand efficiency). Reversing employer/candidate roles generates an alternative stable outcome that favors employees (i.e., an “employee-optimal” outcome that emphasizes supply efficiency). In this case, candidates offer to work for their preferred employer; employers accept their preferred offer and reject unacceptable proposals. As above, the employee-optimal outcome may include unfilled job vacancies and unmatched personnel (i.e., does not clear markets).

Thus, through either emphasis (i.e., on demand or supply efficiency), two-sided matching algorithms offer substantial improvement over hierarchical planning systems (e.g., producing a stable outcome that gives employees a direct role in determining job matches), but they fall short of the efficiencies reflected in market-based labor markets (e.g., unfilled jobs and unemployed candidates imply the markets fail to clear). Similarly, they allow for centralized control over the matching process (e.g., as enabled by hierarchical systems), but they present huge information requirements for effective operation (e.g., as required for efficient labor markets). These represent important considerations in terms of labor-market design. With this background, we now discuss intelligent agent technology and review advances in terms of multi-agent systems developed to enact, automate and support electronic labor markets.

3. Intelligent agent technology

Work in the area of software agents has been ongoing for some time, and it addresses a broad array of applications. Building upon research in the supply chain domain (cf. Mehra and Nissen, 1998; Nissen and Mehra, 1998; Nissen, 2000), agent technology appears to have particular promise to automate and support electronic labor markets. As computational artifacts, they can help overcome human cognitive limitations (e.g., in terms of memory, processing speed), for instance, supporting rapid search and effective filtering through huge numbers of available jobs and potential employees. Further, possessing some artificial intelligence (AI), agents can employ inferential mechanisms (e.g., rules, cases, scripts) to reflect and observe diverse individuals' preferences. In a domain with over a hundred thousand available jobs and prospective employees (e.g., the Navy)—in which both employers and potential employees have specific, idiosyncratic needs and preferences—no other, extant information technology offers the same level of automation and support capability as software agents.

In this section, we discuss some representative, extant agent technologies and then outline key capabilities of the Personnel Mall, a proof-of-concept multi-agent system developed to enact, automate and support an electronic employment market. This agent application emerged through research to adapt the Intelligent Mall—its predecessor application for matching buyers and sellers in products and services markets—to focus on matching people with jobs in labor markets. Our focus on this, implemented, multi-agent system serves to demonstrate the kinds of agent-based electronic employment markets discussed in this article are not simply concepts of theory or items of speculation. Rather, they build upon demonstrated information technology and provide insight into the kinds of electronic marketspaces that are now becoming technically feasible.

3.1. Extant agent applications

Following the literature survey and classification system of Nissen (2000), it is informative to group extant agent applications into four classes: 1) information filtering agents, 2) information retrieval agents, 3) advisory agents, and 4) performative agents. Other groupings from the agents literature (e.g., Bradshaw, 1997; Franklin and Graesser, 1996; Nwana, 1996; White, 1997) could potentially be used as well, but this classification scheme is useful to compare various agent capabilities that are applicable to markets and matching processes. Tables 1–4 are presented to concisely summarize a number of representative, extant agent applications with relevance to markets and matching processes. Beginning with Table 1, briefly, most information filtering agents are focused on tasks such as passively applying user-input preferences to screen and sort e-mail, network news groups, frequently asked questions and arbitrary text.

From Table 2, information retrieval agents address problems associated with collecting information pertaining to commodities such as compact disks and computer equipment, in addition to services such as advertising and insurance. We also include the ubiquitous Web indexing robots in this class along with Web-based agents for report writing, publishing, assisted browsing and other applications listed in the table. Such active information

Table 1. Information filtering agent applications (adapted from Nissen, 2000).

Representative application	Literary source
Filter e-mail messages	Maes, 1994; Malone, Yates, and Benjamin, 1987
Filter network newsgroup postings	Sycara and Zeng, 1996
Filter frequently asked questions (FAQs)	Whitehead, 1994
Filter arbitrary text messages	Verity, 1997

Table 2. Information retrieval agent applications (adapted from Nissen, 2000).

Representative application	Literary source
Collect product/service information	Krulwich, 1996; uVision, 1998; PriceWatch, 1997; Insurance, 1997
Web robots & publication tools	Etzioni and Weld, 1995; Chen et al., 1998; Amulet, 1997; InterAp, 1995
Assisted Web browsing	Burke, Hammond, and Young, 1997
Technical information gathering & delivery	Knobloch and Ambite, 1996; Bradshaw et al., 1997
Shopping "Bots"	Krantz, 1999

Table 3. Advisory agent applications (adapted from Nissen, 2000).

Representative application	Literary source
Recommend compact discs & movies	Maes, 1997; Nguyen and Haddawy, 1999
E-concierge services	Etzioni and Weld, 1995
Campus visit "host"	Zeng and Sycara, 1995
Planning support	Maturana and Norrie, 1997; Pinson, Louca, and Moraitis, 1997
Project coordination advice	Johar, 1997
Computer interface assistance	Ball et al., 1997
Military reconnaissance support	Bui et al., 1996
Financial portfolio advice	Sycara et al., 1996
Buyer/seller matchmaking advice	Freuder and Wallace, 1999
Supply chain decision support	Goodwin et al., 1999

retrieval represents an important market task, and many commercial shopping "bots" are being developed to do this in a variety of product and service markets.

From Table 3, advisory agents are oriented toward providing intelligent advice, and they perform in something of a classical decision-support role. Examples include recommendations for CDs and movies, an electronic concierge, an agent "host" for college campus visits, planning support for manufacturing systems and other applications listed in the table. More recently, work on agents to provide advice on matching buyers with sellers and decision-support agents have addressed the markets and matching domain directly. Unlike commercial shopping "bots" (i.e., information retrieval agents), for instance, agents in this third class decide for themselves which information is needed, plus they seek out and use such information to make recommendations for users.

Finally, from Table 4, performative agents are the most sophisticated with respect to the state of the art, as they often rely on capabilities of agents in other classes (e.g., in-

Table 4. Performative agent applications (adapted from Nissen, 2000).

Representative application	Literary source
Business marketspace	Chavez and Maes, 1996; Fox and Barbuceanu, 2000; Hu, Reeves, and Wong, 1999; Mehra and Nissen, 1998; Nissen and Mehra, 1998; Preece, Hui, and Gray, 1999
Auction marketspace	Rodriguez-Aguilar et al., 1998, Hu, Yen, and Chung, 1999; Sandholm, 1999
Agent negotiation	Bui, 1996; Collins et al., 1998; Guttman, Moukas, and Maes, 1998; Maes, Guttman, and Moukas, 1999; Sun et al., 1999; Tesauro and Kephart, 2000
Scheduling	Sen, 1997; Walsh et al., 1997
Cooperative learning	Boy, 1997
Digital library services	Mullen and Wellman, 1996

formation filtering and retrieval, decision support). But unlike the types of agents from above, performative agents *perform* important knowledge- and information-work activities autonomously, changing the state of the external world (e.g., executing binding commercial transactions) through autonomous, deliberate action. In contrast, the applications from above can only *support* people in their performance of such activities. Examples of performative agents include marketspaces in which agents can conduct business transactions, auction environments in which agents buy and sell on behalf of their users, and several agent system designs for negotiation. Performative agents are also developed to automate knowledge work such as scheduling, autonomously provide a cooperative-learning environment and provide digital library services.

The Personnel Mall is probably best categorized in the fourth group (i.e., performative agents). But it has been designed to build upon agent work in other categories as well. For instance, it exhibits behaviors such as information filtering and retrieval, and it can be used in an advisory role as well as a performative one. Central to the Personnel Mall's potential is its ability to represent a multitude and wide variety of different users—on both the demand and supply sides—to quickly find, retrieve and organize large amounts of market information. And its conformance to market and organizational rules, established for a particular enterprise or circumstance, enables this multi-agent system to automate and support commerce in a broad diversity of electronic markets (e.g., including regulation-laden, hierarchical systems). Such ability suggests the Personnel Mall offers good potential to enact, automate and support the kinds of electronic labor markets addressed through this research.

3.2. The Personnel Mall

As noted above, the Personnel Mall is a proof-of-concept multi-agent system developed to enact, automate and support a Web-based marketspace for employee/job matching. Its ontology and agent behaviors are patterned after the Navy's current, hierarchical planning system. But the design incorporates many aspects of matching algorithms (e.g., explicit consideration of employee/employer preferences) and market-based markets (e.g., open access to market-wide information). In describing this multi-agent system, we first discuss its structure. We then detail the Personnel Mall application and outline its key agent behaviors.

3.2.1. Personnel Mall structure Like its predecessor system—called the "Intelligent Mall," which was developed for matching buyers with vendors in products and services markets, the Personnel Mall employs a shopping mall metaphor for employee/job-matching. In a mall, shoppers are not expected to know in advance which shops exist or what products they offer for sale. Neither are the shops expected to know which other shops are selling like products or with which shoppers they will interact.

Also following its predecessor, the Personnel Mall ontology was established through system design to reflect entities that are common and familiar in corresponding physical markets. For instance, the Intelligent Mall ontology is comprised of entities appropriate for product and service markets (e.g., customers, vendors, products/services, prices). Similarly, the Personnel Mall ontology is comprised of entities appropriate for labor markets (e.g., employees, employers, jobs, wages). But these agent-enabled marketspaces provide a qualitative contrast to their physical marketplace counterparts; that is, instead of *people* searching, matching, buying and selling within physical markets, such market and matching activities are performed instead by *software agents*, agents which serve to represent people (i.e., the users or principals). The agents—implemented through software objects and methods—exchange messages with one another to communicate and coordinate their activities. Mall behaviors (e.g., enacted via methods, rules) are explained in greater detail below.

The mall representation presented in Figure 2 shows intelligent agents representing six sailors, three agents representing commands and one command agent's "shopping list" (i.e., input form) for specifying jobs to be filled through the marketspace agent federation. This mall is truly virtual, in that the sailors do not reside in any single physical location. Indeed, the sailors' physical locations are irrelevant. What is relevant is that an agent is created to represent each sailor interested in participation through this medium. Notice there is no agent or other explicit intermediary to represent the Detailer. Rather, commands' and sailors' agents interact directly through this electronic labor market.

To "launch" a command agent, an officer needs only network access to an agent server (we provide a Web interface for this) and a common Web browser. Any practical number of command agents can co-exist in this mall, and any practical number of (distributed) Mall servers can interlink to define various mall environments of arbitrary size and composition. Thus, the application offers good potential to scale well to large enterprises such as the Navy. Not apparent in the figure is a special agent called "Host." Command agents register their job openings (e.g., on board ships at sea, assigned to ships in port, shore positions), and the Host maintains both a "White Pages" and a "Yellow Pages" directory for participants in the mall. Such directories are used to support both sailors' and commands' search through available job listings and applicants, respectively, and they obviate many human cognitive limitations that would impede market efficiency.

A command agent can be specialized to reflect the preferences and priorities of its principal (e.g., Commanding Officer). For example, a command agent can be specialized to

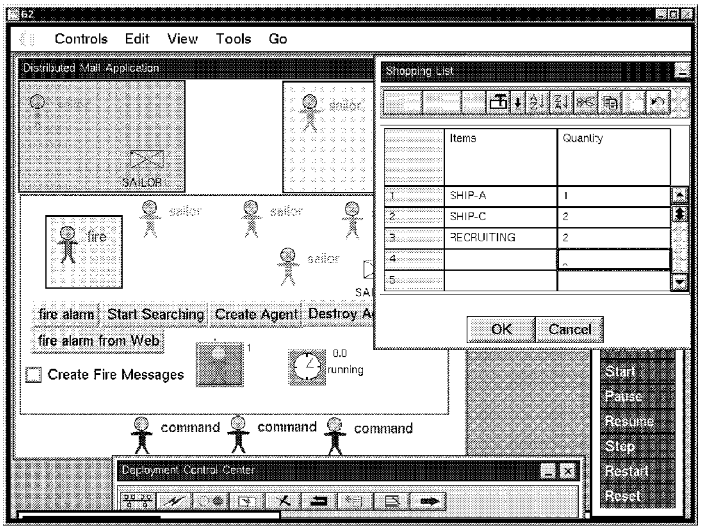

Figure 2. Personnel Mall screenshot.

search and match based on one officer's preferences for education and training over rank and experience, whereas another's (i.e., identical clone) can be specialized to reflect different (even opposing) preferences. And, clearly, each command agent can be instantiated with a unique shopping list of job openings. Other knowledge and information—such as user preferences, budget restrictions, personnel requirements and need dates—can be formalized through rules for the agents.

As an example, say an employer has a job opening. And say this employer prefers employees with technical backgrounds over those with only managerial experience for this job, but he or she refuses to pay more than $2000 per week for such an employee yet needs to hire someone within one week. Using Structured English for explication, two rules to reflect such employer preferences can be written as follows.

Rule 101.
IF applicant-background(X) = technical
AND salary(X) < 2000
AND availability(X) < 1 week
THEN hire(X) = T

Rule 102.
IF applicant-background(X) = managerial
AND salary(X) < 2000
AND availability(X) < 1 week
THEN hire(X) = T

Rules 103–*n*.

In a top-down rule interpreter, Rule 101 is considered first. If all the antecedent conditions of Rule 101 are satisfied by the applicant, then a decision is made to hire this (technical) individual (i.e., "hire = T" or True). Otherwise, Rule 101 does not hire (i.e., a decision is not made to hire this (technical) individual), and Rule 102 is considered in turn. Likewise, if all the antecedent conditions of Rule 102 are satisfied by the applicant, then a decision is made to hire this (managerial) individual. Otherwise, Rule 102 does not hire, the next rule (e.g., Rule 103) is considered in turn, and so forth (e.g., through Rule n). Developing an appropriate set of such rules is accomplished by the knowledge engineer during the agent-design phase, but individual users (e.g., commands) have an ability to modify rules to reflect their particular preferences.

Likewise, sailor agents are designed with a corresponding set of rules, which can be similarly modified by individual users (e.g., sailors) to reflect their particular preferences. Requirements for launching a sailor agent are similar to those pertaining to the commands above. Like their command counterparts, the latter agents represent sailors wishing to participate in the mall and similarly register with the Host. For instance, sailor agents use messages to communicate with the Host and register individual sailors' attributes (e.g., rank, education, training, job skills, availability, community). And sailors can search the Mall to view a complete, current listing of job openings (e.g., using the "White Pages" and "Yellow Pages" services of the Host). Sailors use such listings to identify a specific subset of assignments in which they are particularly interested, and they use an input form to convey their relative preferences among this subset of job assignments.

3.2.2. Job preferences and matching One specific, sailor-agent input form is displayed in Figure 3. Notice the relative preferences listed in the input form for Sailor-2 in this figure. The sailor's first preference is for a job opening on "Ship-C," details (not shown in

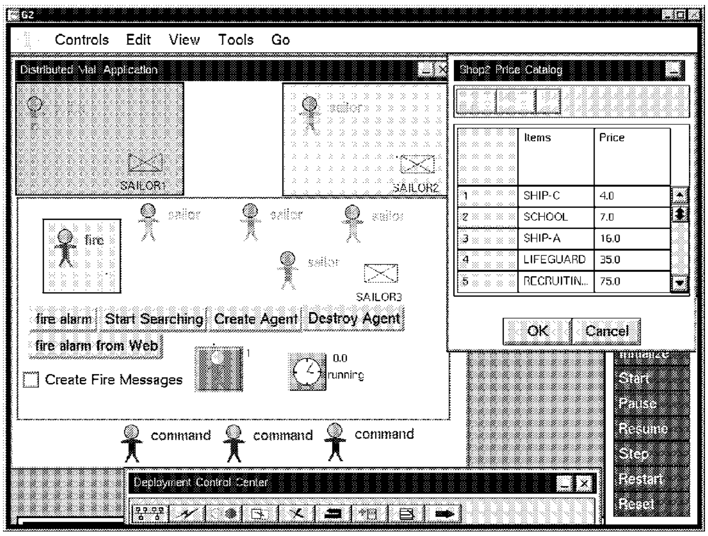

Figure 3. Personnel Mall sailor agent job preferences form.

the figure) of which indicate this assignment involves sea duty onboard a surface combatant scheduled for a tour through a region of potential conflict. Details also indicate this particular assignment provides a good match with the sailor's education, training, skills and experience. Therefore, this, particular assignment appears to represent a good career move for this individual sailor. Clearly, other sailors, as individuals, will have different and unique sets of preferences, which they similarly express through input forms for their own, personal sailor agents.

Continuing down the ranked list of assignment preferences for Sailor-2, this particular sailor would like to attend an advanced training school (e.g., in electronics, computer technology), in the event that he or she is not selected by the Ship-C command. Such advanced training would also represent a good career move for this particular sailor, and the Personnel Mall enables each individual sailor to express such career ambitions and priorities through job rankings. Still lower on the list, this sailor identifies an assignment on Ship-A, which is scheduled to remain in port, as a third-priority alternative. Fourth on the list is a position as lifeguard for a Navy swimming pool in California, and the last item represents a recruiting assignment. Using such sailor-specific information, the Personnel Mall incorporates each sailor's preferences when matching sailors with job openings, and sailors' agents can work to get matches with their highest-priority jobs. With this, the Personnel Mall compares favorably with the outcomes of two-sided matching algorithms discussed above.

The problem with the simple, ranked listings, however, is they do not support relative comparison of preferences *between different sailors*. For instance, Sailor-2 and Sailor-6 may both list the same assignment to Ship-C as their highest-ranked job preference. But whereas Sailor-6 may be desperate for this sea-duty tour—and considers attending the advanced school to be a substantially-inferior alternative—Sailor-2 may only marginally prefer Ship-C duty to the school. If both sailors rank Ship-C and the advanced school as their first- and second-priority jobs, respectively, the Personnel Mall would not have sufficient information to intelligently match the sailor whom most wants each assignment (i.e., assign Sailor-6 to sea duty and Sailor-2 to school). In other words, such ordinal rankings of sailors' job preferences are sufficient to match the most-desired assignments for any individual sailor, but they do not support such comparisons between any two or more sailors. With this, the Personnel Mall again is comparable to the outcomes of the two-sided matching algorithms discussed above

To handle such inter-sailor preferences, the Personnel Mall was enhanced to employ a quasi-price system that enables sailors to *quantify* their relative preferences between alternative job assignments. This effectively transforms the matching problem from one relying on ordinal scales (e.g., ranked job alternatives) to one using ratio scales (e.g., consistent, numerical prices, ratings, scores, credits or other similar quasi-pricing instruments) for quantifying individuals' relative utilities associated with diverse jobs. This is comparable to the manner in which a centralized market (e.g., stock exchange) expresses the relative value of diverse securities in terms of prices per share. This aspect of the Personnel Mall compares favorably with the outcomes of market-based markets discussed above, and expands on the two-sided matching algorithms.

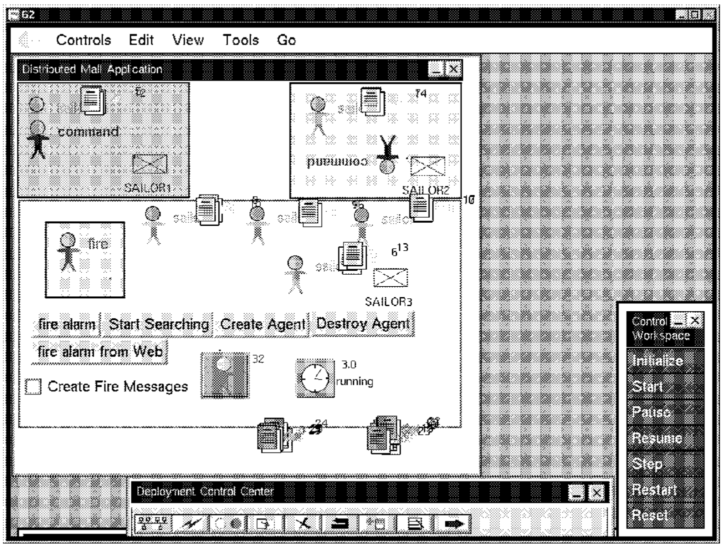

Figure 4. Personnel Mall animation.

3.2.3. Personnel Mall operation Figure 4 presents a still-shot of some animated search and match activities. When a command has specified the jobs to be filled (i.e., using the input form described above), each command agent sends messages to all sailor agents in the Personnel Mall that have expressed interest (i.e., registered with the Host) in jobs on its shopping list. Sailor agents wishing to apply for posted jobs reply directly to the requesting agents and respond to the corresponding messages separately. Each command agent analyzes the messages and determines a preferred sailor for every job on the shopping list, based on quasi-prices.

Notice the command agent visiting Sailor-2 is inverted. Such behavior is explicitly required by "regulation" in this mall environment when sea duty is required, for instance. Although humorous, perhaps, this behavior serves to visibly demonstrate that intelligent agents can be specified to conform to various personnel regulations (e.g., governing sea/shore duty rotations, experiential requirements, Navy needs). Although difficult to see without animation, the other command agent is currently "jumping" up and down, which represents the local custom when interacting with people in the community of Sailor-1. Similar to the regulation-conforming behavior from above, this context-specific behavior serves to demonstrate that intelligent agents can also be sensitive to local customs and variations (e.g., when interacting with sailors from the surface, submarine or aviation communities). The number and kinds of such behaviors that can be specified through intelligent agents are practically unlimited.

When all discussions have been completed with a given sailor agent, the command agent proceeds to the "exit," where messages are exchanged to extend and accept offers, document the match and issue orders. Each command agent then moves to fill the next job on its list, matches the corresponding sailors, and continues in this fashion until all jobs have been filled or no sailors are available to fill jobs still on the list. This latter condition reflects one of excess demand for labor. The Personnel Mall, thus, is subject to the same market principles and dynamics discussed above in terms of labor markets. But the possibility of unfilled high priority jobs and unemployed sailors mirrors a weakness identified in connection with two-sided matching algorithms above. Thus, as with two-sided matching algorithms discussed above, the Personnel Mall appears to offer substantial improvement over hierarchical planning systems (e.g., eliminates the Detailer role, explicitly considers employer and employee preferences) and offers advantages only attainable through electronic markets (e.g., provides access to and helps manage search through abundant, market-wide information, automates many search and matching tasks). But it still falls short of efficiencies enabled through market-based markets (e.g., general equilibrium, market clearing). These represent important considerations in terms of labor-market design.

4. Integrated electronic employment market design

In this penultimate section, we build upon the research and results above to integrate economic and technological elements required for robust electronic employment markets. This design seeks to integrate the best elements of the approaches described above (i.e., two-sided matching algorithms and the Personnel Mall) with additional features required to mimic the properties of market-based labor markets (e.g., general equilibrium). Yet we also seek to preserve some capability for centralized control (e.g., as inherent in hierarchical planning systems). This electronic employment market design, thus, is intended to prescribe requirements for and guide the implementation of the next round of agent-based research and market development. To help guide this discussion, we concentrate on the four, key deficiencies identified to date with the Personnel Mall: 1) unstable outcomes, 2) market non-equilibrium, 3) inadequate control, and 4) complex user interfaces.

4.1. Unstable outcomes

Beginning with the first deficiency, as currently designed, the Personnel Mall is effectively limited to one-sided matching. Specifically, although command agents and sailor agents both reflect and act upon preferences of their principals, the current implementation uses quasi-prices, exclusively, to match sailors with jobs. And because sailors are free to specify their own quasi-prices through sailor agents, job assignments made through the Personnel Mall essentially reflect only the sailors' preferences (i.e., emphasize only supply efficiency). Thus, demand efficiency suffers, and commands are unlikely to be satisfied with job assignments made through the current Mall implementation. Alternatively, the sailors are likely to view this as an improvement over the Navy's current, hierarchical detailing system (i.e., that emphasizes demand efficiency).

As a first step, this suggests incorporating two-sided matching into the Personnel Mall. Drawing from our discussion of two-sided matching algorithms above, by explicitly considering both employer and employee preferences, stable outcomes can be obtained. Although two-sided matching algorithms necessarily continue to emphasize either demand or supply efficiency, at least the outcomes are stable (i.e., neither commands nor sailors would prefer a different, available employee/job match to the one assigned via the algorithm).

This calls for modifying agent behaviors in the Personnel Mall, such that an iterative matching process can be effected and two-sided matching is enforced. Technically, this would require incorporating some kind of intermediary agent—analogous to the centralized organization responsible for two-sided matching (e.g., in the medical residency domain)—conceivably working in conjunction with the Host agent currently implemented in the Mall. The command and sailor agents, in turn, would require modification to interact only through this intermediary—recall command and sailor agents interact directly at present—and sailor/job matches would need to be accomplished via an incremental procedure (i.e., as required by the matching algorithm).

But at least two problems can be identified with this design. First, even after incorporating a capability for two-sided matching, the problems pertaining to two-sided matching algorithms (e.g., market non-equilibrium) will persist. So, although this design would improve upon the current Personnel Mall implementation, it continues to fall short of a market-based market approach. Second, some benefits inherent within the flat architectural design of the Personnel Mall (e.g., sailor and command freedom to interact and match with any available commands and sailors) will be lost, and considerable sailor and command autonomy will be sacrificed (e.g., as matches are relegated to the centralized organization).

Additionally, the Mall's current quasi-pricing scheme requires integration with the two-sided matching algorithm. One approach involves variable bonuses (e.g., quasi-prices), which can be used in a two-sided matching market to reflect job priority. The term *bonuses* suggests that this compensation can take various forms (e.g., wage supplements; benefits, including release time, conference travel; "brownie points" that influence future promotions or job assignments). Bonuses act as compensating wage differentials in competitive labor markets. In particular, bonuses can be offered for the important but undesirable jobs left unfilled. During the iterative matching process, bonuses can change as appropriate until high priority jobs are "voluntarily" filled with appropriately skilled personnel. In the two-sided matching process, this complicates the matching algorithm (Kelso and Crawford, 1982; Roth and Sotomayor, 1990, pp. 171–186). And in the Personnel Mall specification, this would require updating prices before each shopping iteration.

One of the critical issues in this design extension involves how the bonuses are established (e.g., offered by employers, demanded by candidates, or established through a bidding process). Alternative schemes to determine the required bonuses will have implications for how employers and candidates behave (i.e., strategic behavior) and the resulting distribution of labor across jobs. If bonuses are pre-announced and never updated, they may not effectively balance demand and supply for high priority undesirable jobs. If bonuses are systematically increased over time until candidates accept the high priority undesirable jobs, candidates have an incentive to delay the match to receive a higher bonus; but they risk someone else filling the job during their holdout. Finally, bonuses can be determined through a bidding (e.g., auction) process designed to identify the minimum bonus that attracts the candidate most willing to accept the assignment in question (cf. Vickery, 1961; Myerson, 1981, 1983; Milgrom and Weber, 1982). Regardless of the mechanism used to determine the bonus, the impact on the matching process/Personnel Mall is effectively the same. Assignment games provide an alternative approach (cf. Demange and Gale, 1985; Demange, Gale, and Sotomayor, 1986; Roth and Sotomayor, 1990, pp. 202–239; Shubick, 1984, pp. 191–225).

4.2. Market non-equilibrium

As noted above, neither two-sided matching algorithms nor the Personnel Mall can guarantee general equilibrium. And clearly, equilibrium is seldom achieved through hierarchical planning systems. Indeed, both approaches can leave unfilled jobs as well as unemployed sailors, and there is no guarantee that even matched sailors are necessarily the best choices for their particular jobs. Instead, at present, we can only say that the outcome is stable (e.g., through the two-sided matching algorithm) or that inter-sailor preferences are balanced (e.g., through the Personnel Mall). Contrast this with the properties of market-based labor markets, which we note above use dynamic pricing—through pressure to raise or lower relative wages in response to respective labor shortages or surpluses—to ensure partial equilibrium. Moreover, because such wage rates can be compared across industries—as well as jobs—labor is free to move across markets and ensure it is applied to the most-productive jobs/uses (i.e., for general equilibrium).

As a first step, this suggests incorporating a dynamic-pricing mechanism into the Personnel Mall. Drawing from our discussion above, this would involve a mechanism to enable wages (or some equivalent) to adjust automatically in response to shifts in relative labor supply or demand. With this, high priority jobs left unfilled (e.g., after the final round of matching) would increase the relative wage paid or other corresponding benefit (e.g., location, working hours, other amenities) until a qualified employee elected to take it. Conversely, wages/benefits associated with popular jobs would decrease until only one qualified employee elected to take it, or it remained vacant depending on its job priority.

But at least two problems can be identified with this design. First, as noted above, the Navy and other, large organizations represent closed labor markets, in which mid-career labor is very difficult to obtain from the spot market (e.g., outside the current Navy labor pool), and wages are not free to adjust dynamically to relative changes in supply and demand. Rather, wages are relatively fixed and uniform across comparable jobs. So, any dynamic pricing must be accomplished via quasi-prices as opposed to wages. Also, non-pecuniary benefits (e.g., location, working hours, other amenities) tend to be relatively standard in the Navy, so such quasi-prices must be tied to something of value to sailors (and commands) in order to motivate market-like behavior. Specifying the nature and understanding the ramifications of alternative approaches to such quasi-prices represents a challenging research problem.

Second, the mechanism behind price (or quasi-price) dynamics must be specified in a manner that produces equilibrium solutions. This represents a matter for both theoretical and empirical assessment. Such price dynamics also need to be evaluated to ensure they do not exacerbate labor shortages or surpluses (e.g., by incentivizing volatility through speculation) and are reasonably immune to gaming by sailors (and commands). Given continuing problems with achieving such requirements even in well-established, price-based markets (e.g., stock markets), this poses a challenging research problem as well.

4.3. Inadequate control

As noted above, the Personnel Mall empowers sailors and commands to employ intelligent agents and have these software representatives autonomously search the market and find employee/job matches on their behalf. And we note above, the Detailer plays no role in the Mall. However, this design fails to satisfy current Navy expectations regarding control over the personnel assignment process, and such expectations for control are unlikely to be limited to the Navy. Despite the limitations of its current, hierarchical planning system, the Naval Personnel Command has the final say over sailor/job assignments, and its personnel (i.e., detailers) are actively involved in every step of the process (e.g., to ensure demand efficiency). Even the two-sided matching approach offers some level of control (e.g., through a centralized matching organization).

As a first step, this suggests incorporating some control agent (e.g., to represent the Naval Personnel Command) to oversee the process and perhaps even wield veto rights over employee/job assignments. Technically, this approach shares some similarity with incorporating the kind of intermediary agent outlined in the design to address unstable outcomes above. But to exercise control—as opposed to influence—this centralized kind of agent design implies considerably more power and authority than a simple intermediary. And given current limitations in terms of agent technology—not to mention Navy attitudes—it is unclear whether such control agent would even be implemented via software; that is, such control agent(s) may instead represent a role for people to play in the electronic employment market. This requires further study.

Further, as with all multi-agent systems, many properties of the agent *federation* (i.e., collection of Mall agents) as a whole are emergent and difficult to anticipate at the design stage. For instance, two command agents *from the same organization* may compete against one another over the same potential employee (e.g., conceivably driving up wages), even though both agents—and their command principals—would be better off having one agent or the other pursue a different employee candidate. Symmetrically, sailor agents may collude or game the system to produce undesirable outcomes. Many other emergent behaviors (e.g., indecision, deadlock, suboptimization) are also possible, and determining whether the mall metaphor and multi-agent system implementation represent efficient and effective approaches to employer/job matching represents an empirical question for future research.

4.4. Complex user interfaces

Finally, we note above how agents offer good capability to reflect the idiosyncratic preferences of even a large number of diverse users. Indeed, the ability to automate and support a process while attending to so many idiosyncratic user requirements represents a fundamental capability of software agents that sets them apart from other classes of IT. We also note above how individual users can modify their agents (e.g., incorporating rules) to reflect their individual and changing preferences. However, at present, the agent interfaces are inadequate to enable all but the most-sophisticated user (e.g., possessing agent-design experience) to accomplish this activity without technical support. And when one imagines the level of technical support required to assist several hundred thousand such users in an organization the size of the Navy, this represents an impractical approach.

As a first step, this suggests building easy-to-use interfaces to the agents. As an exemplar, recall from the Personnel Mall screenshots above how sailors could complete simple forms to check with the Host's "White Pages" and "Yellow Pages" or specify their job preferences for matching. Similar, simple forms could potentially be developed to enable users to view and modify their agents' searching and matching rules. But with simplicity comes some loss of power, as the range of agent behaviors captured through such simple forms would not compare to that enabled through an ability to build and modify rules directly. This kind of power-versus-simplicity tradeoff is common to all computer-based systems, and striking the appropriate balance for each particular organization (e.g., the Navy) represents a software-design issue. Developing the *capability* for such simple user interfaces, however, represents another important topic for future research.

5. Conclusions and future research

Two modes of matching people with jobs prevail at present: 1) hierarchical planning and 2) distributed markets. Each has strengths and limitations, but few systems have been designed to take advantage of strengths corresponding to both. With evolving information technology, however, the job-matching process could be accomplished far more equitably and efficiently using web-based markets within the firm, and intelligent agents offer excellent potential to help both potential employees and employers find one another in a distributed, electronic marketplace. But realizing this potential goes well beyond simply changing the rules of internal job matching or making agent technology available to job searchers. Rather, the corresponding markets and technologies must be *designed*, together, to mutually accomplish the desired results (e.g., efficient and effective matching) and conform to necessary properties (e.g., market clearing).

Through the research described in this paper, we drew from Game Theory results to assess the feasibility of using two-sided matching algorithms to address this market-design problem. And we drew from current agent research to address the information technology dimension of the problem through implementation of a proof-of-concept multi-agent system to enact, automate and support the corresponding market solution. Working to select and capitalize on the best aspects of two-sided matching and the Personnel Mall approaches to the market-design problem, we integrated the key economic and technological elements required to design robust electronic employment markets. And the corresponding research provides new knowledge and insight into co-development of the requisite economic markets and agent technologies.

However, notwithstanding the new insight developed through this research, a number of problems persist even with our integrated electronic employment market designs. For instance, we noted above how integrating two-sided matching algorithms into the Personnel Mall fails to address deficiencies inherent in such algorithms themselves (e.g., market non-equilibrium), in addition to lost sailor and command autonomy likely to occur in a more-centralized system. As another instance, we noted above how incorporating a dynamic-pricing mechanism into the Personnel Mall fails to consider specific aspects of the closed, Navy labor market (e.g., difficulty obtaining mid-career labor via spot markets, "sticky" wages and benefits). And we indicated that specifying the nature and understanding the ramifications of alternative approaches to dynamic quasi-pricing represents a challenging research problem. These represent two promising topics for future research along the lines of this investigation.

We also identified a need for centralized control over the labor market and outlined how some kind of agent (e.g., human, software) appears necessary to address Navy concerns. Development, implementation and integration of such human or software agents represents another important area for future research, as does continued exploration into the emergent properties and behaviors associated with federations of intelligent (human and software) agents and improving user-interface designs. In addition to such technical research, economic issues—such as how to construct and manage variable bonuses, or how to impede user gaming of electronic employment markets—requires near-term attention.

With this, there is, clearly, much research yet to be accomplished before robust, electronic employment markets are ready for "industrial strength" implementation in the enterprise. And following the lead of this article, such research should necessarily be multidisciplinary and integrative in nature. As noted in the introduction, economic and technological elements associated with labor-market design are inextricably intertwined, and a robust design cannot be achieved without addressing both in an integrated manner. The research described in this paper has taken a first step along these lines. And we hope to help stimulate and guide future work along these lines.

Note

1. It is important to note that all units of labor are homogeneous (interchangeable) within a market-based labor market. Jobs that require different skill levels are considered different labor markets.

References

Amulet. (1997). Amulet online description. `http://www.amulet.com.`

Ball, G., D. Ling, D. Kurlander, J. Miller, D. Pugh, T. Skelly, A. Stankosky, D. Thiel, M. Van Dantzich and T. Wax. (1997). "Lifelike Computer Characters: The Persona Project at Microsoft." In J. Bradshaw (ed.), *Software Agents*. Menlo Park, CA: AAAI Press.

Boy, G.A. (1997). "Software Agents for Cooperative Learning." In J. Bradshaw (ed.). *Software Agents*. Menlo Park, CA: AAAI Press.

Bradshaw, J. (ed.). (1997). *Software Agents*. Cambridge, MA: MIT Press.

Bradshaw, J.M., S. Dutfield, P. Benoit, and J.D. Woolley. (1997). "KAoS: Toward an Industrial-Strength Open Agent Architecture." In J. Bradshaw (ed.), *Software Agents*. Menlo Park, CA: AAAI Press.

Bui, T. (1996). "Intelligent Negotiation Agents for Supporting Internet-based Competitive Procurement." Working paper.

Bui, T., C. Jones, S. Sridar and N. Ludlow. (1996). "Decision Support for Reconnaissance Using Intelligent Software Agents." Naval Postgraduate School research proposal.

Burke, R.D., K.J. Hammond and B.C. Young. (1997). "The FindMe Approach to Assisted Browsing." *IEEE Expert*, July/August, 32–40.

Chavez, A. and P. Maes. (April 1996). "Kasbah: An Agent Marketplace for Buying and Selling Goods." In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, London, UK.

Chen, H., Y. Chung, M. Ramsey and C. Yang. (1998). "An Intelligent Personal Spider (Agent) for Dynamic Internet/Intranet Searching." *Decision Support Systems* 23, 41–58.

Collins, J., B. Youngdahl, Sc. Jamison, B. Mobasher, and M. Gini. (1998). "A Market Architecture for Multi-Agent Contracting." In K. Sycara and M. Wooldridge (eds.), *Proceedings of the Second International Conference on Autonomous Agents*, Minneapolis, MN, pp. 285–292.

Crawford, V.P. (1988). "Comparative Statics in Matching Markets." *Journal of Economic Theory* 54, 389–400.

Demange, G. and D. Gale. (1985). "The Strategy Structure of Two-Sided Matching Games." *Econometrica* 53, 873–888.

Demange, G., D. Gale, and M. Sotomayor. (1986). "Multi-Item Auctions." *Journal of Political Economy* 94, 863–872.

Ehrenberg, R.G. and R.S. Smith. (1997). *Modern Labor Economics: Theory and Public Policy*. Reading, MA: Addison-Wesley.

Etzioni, O. and D.S. Weld. (1995). "Intelligent Agents on the Internet: Fact, Fiction, and Forecast." *IEEE Expert*, August, 44–49.

Fox, M. and M. Barbuceanu. (2000). "The Integrated Supply Chain Management Project." Enterprise Integration Laboratory, Department of Industrial Engineering, University of Toronto. `http://www.eil.utoronto.ca/iscm-descr.html`.

Franklin, S. and A. Graesser. (1996). "Is It an Agent or Just a Program? A Taxonomy for Autonomous Agents." In *Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages*. New York, NY: Springer-Verlag.

Freuder, E.C. and R.J. Wallace. (1999). "Matchmaker Agents for Electronic Commerce." In: T. Finin and B. Grosof (eds.), *Proceedings American Association for Artificial Intelligence*. Orlando, FL, Workshop on AI for Electronic Commerce, July.

Gale, D. and L. Shapley. (1962). "College Admissions and the Stability of Marriage." *American Mathematical Monthly* 69, 9–15.

Goodwin, R., P. Keskinocak, S. Murthy, F. Wu, and R. Akkiraju. (1999). "Intelligent Decision Support for the e-Supply Chain." In T. Finin and B. Grosof (eds.), *Proceedings American Association for Artificial Intelligence*. Orlando, FL, Workshop on AI for Electronic Commerce, July.

Guttman, R.H., A.G. Moukas, and P. Maes. (1998). "Agent-Mediated Electronic Commerce: A Survey." *Knowledge Engineering Review*, June. `http://ecommerce.media.mit.edu/`.

Harrison, G.W. and K.A. McCabe. (1989). "Stability and Preference Distortion in Resource Matching: An Experimental Study of the Marriage Market." Department of Economics, University of New Mexico, mimeo.

Hu, J., D. Reeves, and H.S. Wong. (1999). "Agent Service for Online Auctions." In T. Finin and B. Grosof (eds.), *Proceedings American Association for Artificial Intelligence*. Orlando, FL, Workshop on AI for Electronic Commerce, July.

Hu, J., J. Yen, and A. Chung. (1999). "A Virtual Property Agency: Electronic Market with Support of Negotiation." In T. Finin and B. Grosof (eds.), *Proceedings American Association for Artificial Intelligence*. Orlando, FL, Workshop on AI for Electronic Commerce, July.

Insurance. (1997). Insurance online description. `http://www.dmatters.co.uk`.

InterAp. (1995). "InterAp Assigns Intelligent Agents to the Web." *PCWeek*, 12 June.

Johar, H.V. (1997). "SoftCord: an Intelligent Agent for Coordination in Software Development Projects." *Decision Support Systems* 20, 65–81.

Kelso, A.S., Jr. and V.P. Crawford. (1982). "Job Matching, Coalition Formation and Gross Substitutes." *Econometrica* 50, 1483–1504.

Knobloch, C.A. and J.L. Ambite. (1997). "Agents for Information Gathering." In: J. Bradshaw (ed.), *Software Agents*. Menlo Park, CA: AAAI Press.

Krantz, M. (1999) "The Next E-volution: BusinessBots Could Transform Corporate Commerce, Just as the Web Transformed Consumer Shopping." *TIME*, 12 July, 47.

Krulwich, D. (1996). *An Agent of Change*. Andersen Consulting Center for Strategic Technology Research.

Maes, P. (1994). “Agents that Reduce Work and Information Overload.” *Communications of the ACM* 37(7), 30–40.

Maes, P. (1997). “Pattie Maes on Software Agents: Humanizing the Global Computer.” *Internet Computing*, July–August.

Maes, P., R.H. Guttman, and A.G. Moukas. (1999). “Agents that Buy and Sell.” *Communications of the ACM* 42(3), 81–87.

Malone, T.W., J. Yates, and R.I. Benjamin. (1987). “Electronic Markets and Electronic Hierarchies.” *Communications of the ACM* 30(6), 484–497.

Maturana, F.P. and D.H. Norrie. (1997). “Distributed Decision-making Using the Contract Net Within a Mediator Architecture.” *Decision Support Systems* 20, 53–64.

Mehra, A. and M.E. Nissen. (1998). “Case Study: Intelligent Software Supply Chain Agents Using ADE.” In *Proceedings from the AAAI Workshop on Software Tools for Developing Agents*.

Milgrom, P.R. and R.J. Weber. (1982). “A Theory of Auctions and Competitive Bidding.” *Econometrica* 50, 1089–1122.

Mongell, S.J. and A.E. Roth. (June 1991). “Sorority Rush as a Two-Sided Matching Mechanism.” *American Economic Review* 81, 441–464.

Moore, M.J. and W.K. Viscusi. (1990). *Compensation Mechanisms for Job Risks: Wages, Worker’s Compensation, and Product Liability*. Princeton, NJ: Princeton University Press.

Mullen, T. and M.P. Wellman. (November 1996). “Market-Based Negotiation for Digital Library Services.” In *Second USENIX Workshop on Electronic Commerce*.

Myerson, R.B. (1983). “The Basic Theory of Optimal Auctions.” In R. Englebrecht-Wiggins, M. Shubik and R. Stark (eds.), *Auctions, Bidding and Contracting: Uses and Theory*. New York: New York University Press.

Myerson, R.B. (1981). “Optimal Auction Design,“ *Mathematics of Operations Research* 6, 58–73.

Nguyen, H. and P. Haddawy. (1999). “DIVA: Applying Decision Theory to Collaborative Filtering.” In T. Finin and B. Grosof (eds.), *Proceedings American Association for Artificial Intelligence*. Orlando, FL, Workshop on AI for Electronic Commerce, July.

Nissen, M.E. and A. Mehra. (1998). “Redesigning Software Procurement through Intelligent Agents.” In *Proceedings of the AAAI Workshop on AI in Reengineering and Knowledge Management*.

Nissen, M.E. (2000). “Agent-Based Supply Chain Integration.” *Journal of Information Technology Management*, forthcoming.

Nwana, H.S. (1996). “Software Agents: An Overview.” *Knowledge Engineering Review* 11(3), 205–244.

Pinson, S., J.A. Louca, and P. Moraitis. (1997). “A Distributed Decision Support System for Strategic Planning.” *Decision Support Systems* 20, 35–51.

Preece, A., K. Hui, and P. Gray. (1999). “KRAFT: Supporting Virtual Organisations through Knowledge Fusion.” In T. Finin and B. Grosof (eds.), *Proceedings American Association for Artificial Intelligence*. Orlando, FL, Workshop on AI for Electronic Commerce, July.

PriceWatch. (1997). PriceWatch online description. `http://www.pricewatch.com.`

Rodriguez-Aguilar, J.A., F.J. Martin, P. Noriega, P. Garcia, and C. Sierra. (1998). “Competitive Scenarios for Heterogeneous Trading Agents.” In K. Sycara and M. Wooldridge (eds.), *Proceedings of the Second International Conference on Autonomous Agents*. Minneapolis, MN, pp. 293–300.

Roth, A.E. (1991). “A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians in the U.K.” *American Economic Review* 81 June, 415–440.

Roth, A.E. (1984). “The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory.” *Journal of Political Economy* 92, 991–1016.

Roth, A.E. and E. Peranson. (1997). “The Effects of Change in the NRMP Matching Algorithm.” *The Journal of the American Medical Association* 278, 729–733.

Roth, A.E. and M.A.O. Sotomayor. (1990). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge, UK: Cambridge University Press.

Sandholm, T. (1999). “eMediator: A Next Generation Electronic Commerce Server.” In T. Finin and B. Grosof (eds.), *Proceedings American Association for Artificial Intelligence*. Orlando, FL, Workshop on AI for Electronic Commerce, July.

Sen, S. (1997). “Developing an Automated Distributed Meeting Scheduler.” *IEEE Expert*, July/August, 41–45.

Shubick, M. (1984). *A Game Theoretic Approach to Political Economy*. Cambridge, MA: MIT Press.

Sun, R., B.T. Chu, R. Wilhelm, and J. Yao. (1999). "A CSP-Based Model for Integrated Supply Chains." In T. Finin and B. Grosof (eds.), *Proceedings American Association for Artificial Intelligence*. Orlando, FL, Workshop on AI for Electronic Commerce, July.

Sycara, K., A. Pannu, M. Williamson and D. Zeng. (1996). "Distributed Intelligent Agents." *IEEE Expert*, December, 36–46.

Sycara, K. and D. Zeng. (1996). "Coordination of Multiple Intelligent Software Agents." *International Journal of Cooperative Information Systems*.

Tesauro, G.J. and J.O. Kephart. (2000). "Foresight-Based Pricing Algorithms in an Economy of Software Agents." IAC Reports, IBM, `http://www.ibm.com/iac/papers/ice98_fs/fs_public.html`.

U.S. Navy. (2000). Bureau of Naval Personnel Homepage, `www.persnet.navy.mil`.

U.S. Navy. (undated). Commander in Chief, United States Pacific Fleet. "Enlisted Distribution System Overview." Power Point Presentation.

uVision. (1998). `http://www.uvision.com/`.

Verity. (undated). Verity online description, `http://www.verity.com`.

Vickery, W. (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* 16, 8–37.

Walsh, W.E., M.P. Wellman, P.R. Wurman, and J.K. MacKie-Mason. (1997). "Some Economics of Market-Based Distributed Scheduling." In *Proceedings International Conference on Multi-Agent Systems*.

White, J. (1997). "Mobile Agents." In J. Bradshaw (ed.), *Software Agents*. Menlo Park, CA: AAAI Press.

Whitehead, S.D. (1994). "Auto-faq: An Experiment in Cyberspace Leveraging." In *Proceedings of the Second International WWW Conference*, Vol. 1, pp. 25–38.

Zeng, D. and Sycara, K. (1995). "Cooperative Intelligent Software Agents." Carnegie Mellon University Technical Report no. CMU-RI-TR-95-14, March.

Mark Nissen is an Assistant Professor of Information Systems and Acquisition Management at the Naval Postgraduate School. His research focuses on the investigation of knowledge systems for enabling and managing change in areas such as process innovation, electronic business and knowledge flow. Recently he has been investigating knowledge systems to innovate processes in the acquisition domain, and he is currently involved with intelligent supply chain agents, as well as techniques and technologies for the capture and distribution of knowledge in very-large enterprises. Mark's publications span both the information systems and acquisition fields, with recent and forthcoming articles in journals such as *MIS Quarterly*, *Journal of Management Information* Systems, *Decision Support Systems*, *Journal of Information Technology Management*, *Acquisition Review Quarterly* and *National Contract Management Journal*. He has also recently published his first book, entitled *Contracting Process Innovation*, and he serves as Program Manager for the External Acquisition Research Program, which is sponsored by the Defense Acquisition University. Before his information systems doctoral work at the University of Southern California, he acquired over a dozen years' management experience in the aerospace and electronics industry and served as a Supply Officer in the Naval Reserve.

William R. Gates is an Associate Professor of Economics at the Naval Postgraduate School. He received his PhD in economics from Yale University. In addition to his work in game theory and assignment problems, his research interests include incentives and asymmetric information, defense economics and the economics of international defense alliances. His research has been published in *International Studies Quarterly*, *International Public Management Journal*, *Defense Analysis* and *The Journal of Cost Analysis*. Before joining the Naval Postgraduate School, Professor Gates was an economist at NASA's Jet Propulsion Laboratory.