

Procedures for Updating Classification Systems: A Study of Biotechnology and the Standard Occupational Classification System

Neil Malhotra¹ and Jon A. Krosnick²

Government, business, and academic statistical organizations routinely develop classification systems to compartmentalize occupations, fields of study, foods, and many more classes of objects. When innovations occur in living and working conditions, so must innovations occur in these classification systems. This article explores the feasibility of applying cognitive psychology research techniques as a tool to guide such updating. The method entails two cognitive exercises and three analytical approaches that assist experts in identifying the deficiencies in an existing classification system. To illustrate application of the procedure, the method was applied to the Standard Occupational Classification System (SOC) in an effort to accommodate recent changes in the biotechnology industry. Various indicators attest to the validity of the results and therefore encourage use of the methodology with other classification systems and innovations.

Key words: Classification systems; SOC; employment; sorting; cluster analysis; biotechnology.

1. Introduction

One of the most common statistical tasks of businesses and government research agencies is classification: the organization of units into hierarchies of categories. For instance, the U.S. Census Bureau classifies industries within the North American Industry Classification System (NAICS) for the purposes of calculating standardized business statistics; the National Science Foundation's Science Resources Statistics Program organizes academic fields of study when tracking career development in science and engineering fields; and the Department of Health and Human Services categorizes Activities of Daily Living (ADL), such as eating and dressing, to measure the prevalence of disability in the population.

¹ Stanford University, Department of Political Science, Encina Hall West, Room 100, Stanford, CA 94305-6044, U.S.A. Email: neilm@stanford.edu

² Stanford University, Departments of Communication, Political Science, and Psychology, 434 McClatchy Hall, 450 Serra Mall, Stanford, CA 94305, U.S.A. Email: krosnick@stanford.edu

Acknowledgments: This article is based on a study that was commissioned by the U.S. Bureau of Labor Statistics (BLS) in collaboration with the Center for Human Resource Research at Ohio State University and the National Opinion Research Center (NORC) at the University of Chicago. The authors thank Michael McElroy, Michael Horrigan, Michael Pilot, Norman Saunders, Theresa Cosca, Bernard Dugoni, Micah Sjoblom, Jake Bartolone, Gina Wellmaker, and Alexander Tahk for their help and advice. Jon Krosnick is University Fellow at Resources for the Future. Detailed statistics are available from the authors upon request.

One especially important focus of classification efforts for research is occupations. For decades, U.S. federal statistical agencies have relied heavily upon descriptions of the American workforce made using such classification systems. During the last decade, a consortium of U.S. federal agencies has devoted considerable resources to improving these systems. Their primary goal has been to develop and maintain a classification system that encompasses the totality of the work being done in America and also is sufficiently flexible to be useful to a diverse group of agencies, including the U.S. Bureau of Labor Statistics (BLS), the U.S. Census Bureau and the Employment and Training Administration.

The result of the last revision process – the 1998 Standard Occupational Classification System (SOC) – is a four-tiered structure that includes 23 major occupational groups; within them, in a hierarchical arrangement, are 98 minor groups, 452 broad occupations, and 822 detailed occupations. The initial work for the next round of SOC revision has already commenced, and release of an updated system is expected in time for the 2010 Census of the Population. No doubt, an important goal of this work is to produce a revision that accurately and completely updates the current SOC to reflect the considerable occupational changes that have occurred in the nation since 1998.

In this article, we explore the feasibility of applying research techniques derived from principles of cognitive psychology to updating and expanding classification systems, via a case study of the Standard Occupational Classification System. We set out to develop methodologies for identifying new and emerging occupations and for determining whether those occupations were already adequately represented in the SOC or whether amendments to the SOC might help to accommodate them. In addition, we explored whether methodologies could be developed to identify occupational categories that should be added to the SOC, regardless of whether they reflect new and emerging occupations or not. As we discuss below, these procedures are exportable to the revision of other classification schemes, such as the NAICS and the ADL.

Our focus in this study was the field of biotechnology, chosen because it is fast-developing and because industry experts have provided feedback to BLS indicating that the SOC should be expanded to accommodate biotechnology's new and emerging occupations. However, our goal was more general: to develop and test a set of procedures for identifying needed additions to any classification system, the SOC and others alike.

We begin below by presenting a very brief history of the development and uses of the SOC. The following section provides an overview of our methodology for updating classification systems. The fourth and fifth sections explicate the data collection procedures and analytical techniques that comprise our method. A case study of biotechnology and the SOC is used to illustrate the method. The sixth section summarizes the findings and provides recommendations for researchers wishing to use our procedures to improve other classification schemes.

2. Development and Uses of the SOC

The development of the SOC has been described extensively elsewhere (Levine et al. 1999; Salmon 1999; Pollack et al. 2002). Its introduction in 1977 reflected the need for a standard system of classifying occupations that could be used by a diverse group of U.S.

federal agencies. Previously, various agencies had developed their own classification systems, and the U.S. Employment Service constructed a patchwork system to link these separate systems together. The SOC was revised in 1980, but no significant changes were made until an extensive revision process commenced in 1994 (Levine et al. 1999).

Critics argued that the SOC was out-of-date and did not adequately reflect changes that had occurred in rapidly evolving fields, such as production and technology (Kelleher et al. 1993). A central goal of the revision process was to construct a system that focused on the *work* performed (as opposed to the industry in which a job was located), as well as on the skills and education needed to perform that work. In 1998, the product of this effort (which involved more than a dozen agencies) was released as the current SOC (Levine et al. 1999). In 2005, a new effort was launched to update the SOC in time for the 2010 Census.

Occupational data are widely used by researchers in government, business, and academia, underscoring the need for a comprehensive, up-to-date, and flexible system. U.S. federal statistical agencies utilize the SOC in surveys of businesses to assess which areas of the economy are growing and which are stagnating. These data inform education, job training, and immigration policies, among others. Employers depend on the SOC to set salary scales, make industry comparisons, and plan workforce and facility expansions (U.S. Bureau of Labor Statistics 2004). Government surveys ask businesses for open-ended reports of the number of employees within each SOC category. Open-ended questions are appealing, because they allow respondents to give whatever answers they wish, without being drawn to particular ones as in the case of closed-ended questions (e.g., Schuman and Presser 1996). Categorizing responses to open-ended questions for statistical analysis requires a comprehensive and clear classification system, though.

3. Overview of Methodology

The cognitive exercises we designed were intended to (1) measure professional experts' beliefs about which occupations are new and emerging in a field, (2) gauge the compatibility of all occupations in a profession with the SOC, (3) identify occupations that are not optimally represented in the current SOC, and (4) identify new categories that could be added to the SOC to optimally accommodate these occupations.

The philosophy underlying our approach is as follows: We presume first that the changing nature of work, especially in fields using new technologies, can yield the creation of new jobs, the activities involved in which do not necessarily fit well into existing SOC categories. We presume that experts who are very familiar with the types of work being done in the field are well-equipped to assess the goodness of fit of jobs to the existing SOC and that the mental representations of jobs that these experts hold in their long-term memories reveal similarities and differences between those jobs and can be used to detect an optimal structure of them. However, we adopted skepticism typical of contemporary psychology, assuming that these experts might not necessarily be able to accurately identify jobs that do not fit the SOC if asked directly to do so. Rather, the best measurements of poor fit might be what psychologists call "implicit measures" (e.g., Schacter 1987), ones that involve observing people's cognitive activities to detect signs of poor fit of jobs to categories, without asking people directly to report this fit.

To this end, we designed two procedures: (1) *an assignment task*, in which participants took a comprehensive set of jobs and assigned each job to a SOC category and made a series of judgments about each assignment (these participants also identified the jobs they believed were new and emerging in the field), and (2) *a sorting task*, in which participants sorted the comprehensive set of jobs into categories representing their commonalities. These tasks permitted analysis of the resulting data using indirect indicators to identify instances of poor fit and comparison of them to explicit reports of fit provided by research participants.

To test the feasibility of these procedures, we recruited 22 human resources executives from biotechnology companies in the San Francisco Bay Area to complete the cognitive exercises (see Appendix 1 for details on the recruitment procedure). Thirteen performed the assignment task, and nine performed the sorting task. Because this was a pilot study, we were only permitted to present each item to nine participants. The U.S. Office of Management and Budget (OMB) requires clearance for any study with more than nine subjects. Because clearance can often take up to six months, time pressures required us to limit the number of participants. During the assignment task, each participant evaluated as many descriptions as he or she could in the allotted two-hour time period. Five participants completed evaluations of all 172 job descriptions, and four did not. Consequently, four additional participants were recruited to evaluate the descriptions that these individuals did not assess.

We then analyzed the resulting data using a variety of statistical techniques to identify new and emerging occupations in biotechnology and to identify categories that might be added to the SOC to optimally accommodate them. Most importantly, we compared the results of various assessment methods to gauge their reliability and validity. In this instance, validity is defined as the extent to which the method accurately taps the beliefs of industry experts regarding needed additions to the SOC. Because the federal government has long relied on such experts to guide revision of the SOC, reliance on them in the present research context seems well-justified.

The procedures we implemented required stimuli in the form of a set of job titles and descriptions that encompassed all of the work being done in biotechnology. This information was provided to us by Radford Surveys, a compensation and benefits market research firm (based in San Jose, California) that conducts human resources surveys for corporations to use in directing their hiring practices. Each year since 1985, Radford has conducted the Radford Biotechnology Survey, which currently collects data from more than 500 biotechnology firms across the U.S.

A centerpiece of the Radford survey is its listing of all jobs being performed in the field, which numbered 172 when our study was conducted. Because these job titles and descriptions are proprietary and confidential, we cannot report them in this article. Radford constantly updates this list based on information obtained from their participating firms. Both Radford and its survey respondents have incentives to ensure that job descriptions are accurate and comprehensive and up-to-date. When a respondent believes that a job is missing from Radford's list, the respondent informs Radford, and Radford adds the job. The biotechnology firms take great care in updating this job listing, because they receive and depend upon the survey results. Only by carefully reporting their human resources data can the companies cooperatively produce a useful set of survey results. Consequently, the Radford Biotechnology Survey mobilizes firms to provide a collective public good in the form of industry intelligence, which biotechnology firms use in determining their

human resources practices. The Radford job list therefore seems likely to be comprehensive and accurate.

Applying our procedures to other industries and classification systems would require obtaining similar stimuli. However, it would not be necessary to rely upon a Radford-like firm. If one wanted to focus on occupations, a sample of firms in an industry could be contacted (using a member directory of that industry’s professional association), and these firms could be asked to provide open-ended paragraph-long descriptions of the work being done by their employees. Alternatively, focus groups could be conducted with industry experts to generate lists of jobs and work tasks.

Three sets of statistical analyses were conducted using data from the assignment task and the sorting task to identify inadequacies in classification systems. First, we used data from the assignment task to identify occupations that were incompatible with the existing SOC. We then grouped these individual occupations into more general categories using the pile label data from the sorting task. Second, we repeated this procedure for occupations respondents identified as “new and emerging” in the assignment task. Third, we performed a cluster analysis of the sorting task data to assess the cognitive structure of biotechnology occupations, again using the pile labels to organize and label the clusters. As described below, each of these techniques entailed both quantitative and qualitative analysis, necessitating some familiarity with the domain. We then qualitatively amalgamated results produced by the three analytical methods to identify the categories missing from the SOC.

A flowchart summarizing the four stages of the methodology (identification of stimuli and subjects, data collection procedures, data analysis, and summarization) is presented in Figure 1.

4. Data Collection Procedures

In this section, we describe the data collection procedures for the two cognitive exercises:

(1) the SOC assignment task; and (2) the sorting task. As illustrated in Figure 1, the recruited biotechnology experts completed the two tasks, which use the 172 job descriptions as stimuli.

4.1. SOC Assignment Task

The first cognitive exercise we designed asked participants to read each of the 172 biotechnology job descriptions and assign them to the SOC category that fit best.

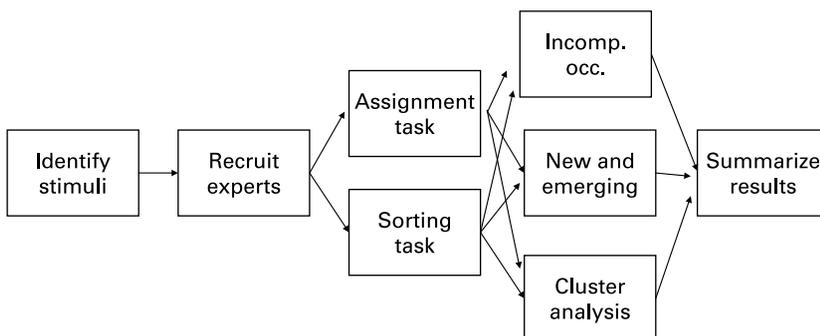


Fig. 1. Flowchart of general methodology

A computer presented one job description at a time on the screen in an order uniquely randomized for each participant.

To minimize wasted time, we eliminated 381 SOC categories that clearly had no relevance to biotechnology (e.g., “Entertainers and Performers,” “Sports and Related Workers”), reducing the available SOC codes from 822 to 441. These SOC categories were printed on one set of pieces of 8.5" × 11" paper that were pinned to large pieces of poster board next to the computer screen for easy viewing and on another set of loose pieces of paper that were handed to participants. Before beginning the exercise, participants were given time to familiarize themselves with the SOC (see Appendix 2 for the instructions given to the participants). We did not record the amount of time each participant took to familiarize herself with the SOC. In future research, we recommend that training time be standardized across subjects.

Five measurements were made for each job description. Four of these measurements were designed to assess the goodness of fit of the job description to an SOC category. First, we noted the SOC category to which the job description was assigned. Additionally, the computer recorded the number of microseconds between the appearance of the job description on the screen and the participant's entering the SOC code into the computer.³ The longer it took a participant to assign a job description, the greater was the potential that the SOC did not offer a suitable category to represent the job. Beginning with pioneering studies conducted by Donders (1969), reaction time has been a standard objective measure used in cognitive psychology research to gauge processing ability. Numerous studies have found that reaction time increases with task difficulty (Welford 1980). Consequently, the less compatible a job is with the SOC, the harder it should be to select a matching SOC code, and the longer it should take to make the selection.

Although we could not measure the time it took to implement each cognitive step involved with the categorization process separately, we did attempt to identify the sources of slow decision-making. A slow reaction time could occur for either of two reasons: (1) the participant could have failed to find a suitable SOC category and spent time choosing among suboptimal categories, or (2) the SOC could have offered multiple suitable categories for a single job, so the participant spent time choosing among a set of plausible competitors. To decompose the slow reaction times into these two sources, we asked participants two additional questions after they made each SOC assignment:

- “How well does this statement fit the category you assigned it to?” (response choices: “not at all,” “slightly well,” “moderately well,” “very well,” “extremely well”).
- “How confident are you that you have assigned this statement to the best-fitting category in the set?” (response choices: “not at all,” “slightly confident,” “moderately confident,” “very confident,” “extremely confident”).

³ We do not presume that distinctions in terms of microseconds are psychologically meaningful and do not interpret our results as such. We simply took advantage of the computer's power to achieve such precision so as to minimize measurement error as much as possible.

The first of these questions was designed to identify jobs that did not fit any SOC category, and the second question was designed to identify jobs that fit multiple SOC categories well. This is not to say that these two constructs should not be highly correlated; indeed, as explained below, jobs that fit well also inspired greater confidence. However, there is some uniqueness to the two judgments. Other factors could have also influenced reaction time (e.g., length of the description, respondent experience and fatigue, differences in the baseline speeds), which we discuss in further detail below.

We also sought to identify jobs that were new and emerging in the field by asking participants a third question after they made each SOC assignment:

- “Would you say that this is a relatively new job in biotechnology that will be increasingly common in the future?” (response choices: “new and increasing,” “*not* new and increasing”).

It is possible that a job could be new and *not* increasing (and vice versa). However, a new occupation that is not becoming a more integral part of the field should not be a high-priority candidate for revision of the classification system. And an old occupation that is increasing is presumably already reflected in the SOC. Nonetheless, it is likely that “new” and “increasing” are highly correlated attributes, because recently developed work approaches are likely to drive the field as it develops.

Thus, we obtained five data points per participant per job description ($9 \times 172 = 1,548$ observations total): SOC assignment, reaction time, fit, confidence, and whether the job is new and emerging.

Using both explicit assessments and implicit measures such as reaction time, studies by Rosch (1975a, 1975b), Rosch et al. (1976) and Rips et al. (1973) on the psychology of categorization found that people prefer basic level categories (i.e., those from which people can form concrete images) to their superordinate categories (i.e., collections of basic categories) and subordinate categories (i.e., divisions of basic categories). People also prefer typical instances of categories to atypical ones. Accordingly, classification systems in which the categories have a psychological reality for domain experts may exhibit greater utility. In the context of the SOC, a psychologically optimal system would allow coders to accurately classify open-ended descriptions by survey respondents as well as facilitate the modification of categories as expert opinion evolves. The explicit and implicit measures elicited by the SOC assignment task permit the creation of new categories that closely fit the thinking of experts in the biotechnology domain and, according to the psychological literature, should bring benefits with respect to efficiency and efficacy.

4.2. *Sorting Task*

With the sorting task, we sought to assess the categories that industry experts generate naturally to organize the work being done in biotechnology. Participants were presented with a stack of 172 index cards with a Radford job description printed on each one. The cards were presented to each of the nine participants in unique random order. Participants then sorted the cards into piles, creating a different stack for each different category of jobs. Participants wrote a label for each pile on a blank piece of paper to describe the types

of jobs in the pile (see Appendix 3 for the instructions provided to participants for the sorting task). Sorting tasks have been widely used in cognitive and social psychology to assess cognitive structure. For methodological overviews, see Brewer and Lui (1989a) and Coxon (1999). For other applications, see Brewer and Lui (1989b) and Chi et al. (1981).

5. Analytical Methods

As illustrated in the center of Figure 1, the three analytical methods we used combine data from the two cognitive exercises in order to identify inadequacies in a classification system. To explicate the methods, we describe in detail the case study of revising the SOC to accommodate recent changes in the biotechnology industry.

5.1. Analysis 1: Identifying All Biotechnology Occupations Not Represented in the SOC

The first method is to identify all items not well represented by the existing classification scheme. The steps of the method are presented in Figure 2. In our case, we gauged the fit of the 172 biotechnology jobs to the SOC with the goal of identifying jobs that are not represented there. To do so, we employed five indicators of goodness of fit collected during the SOC assignment task for each job description:

- (a) *Agreement*, the largest percent of participants who assigned the description to the same SOC code (i.e., the size of the majority or the greatest plurality). In some cases, there was a tie among pluralities. For instance, if 4 out of 9 respondents assigned a job to the “Biological Scientists, All Other” SOC category, 4 assigned it to the “Life Scientists, All Other” category, and 1 assigned it to the “Scientists, Research” category, then the agreement score would be 44.4% (4/9). Alternatively, if 7 out of 9 respondents assigned a job to the “Biological Scientists, All Other” category, then the agreement score would be 77.8% (7/9). A high level of agreement indicates that the SOC adequately represents a given job description.
- (b) *Fit*, indicating how well participants said the selected SOC code fit the job description on a 1 to 5 scale, 1 indicating that the SOC code fits “not at all,” and 5 indicating that

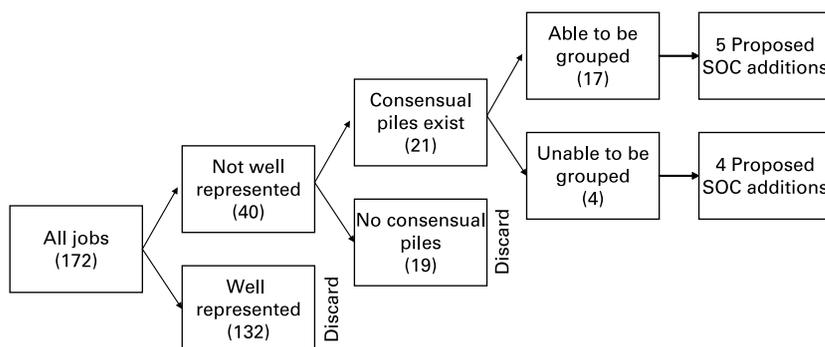


Fig. 2. Flowchart of identification of incompatible occupations

- the SOC code fits “extremely well”. A higher fit score represents a higher level of compatibility with the SOC.
- (c) *Confidence*, indicating the participant’s certainty that he or she selected the best-fitting SOC code from among all SOC codes. These ratings were made on a 1 to 5 scale, with larger numbers indicating greater confidence. Greater confidence was taken to indicate greater compatibility of the job with the SOC.
- (d) *Reaction time*, indicating how quickly participants assigned an SOC code to a job description. We calculated these times adjusting for the length of the job description (because longer descriptions would presumably take longer to read), the number of job descriptions that the participant had read before working on this one (because more practice with the assignment task would presumably make people faster at it), and baseline differences between participants in how quickly or slowly they made all assignments (because some people are generally faster at this task than others, and different people rated different job descriptions). See Appendix 4 for an explanation of the calculation procedure and a regression documenting the validity of the reaction time measurements. Faster reaction times are a sign of a higher level of compatibility of a job with the SOC.
- (e) A standardized *composite* of the above four measures. To create the composite score, reaction times were recoded such that larger numbers indicated faster times. For each measure for each job description, we subtracted the mean and divided by the standard deviation and then summed the standardized scores to yield a single composite score. Larger numbers indicated better correspondence of a job to the SOC.

Table 1 shows that the first four indicators of goodness of fit were all moderately to strongly positively correlated with one another in expected directions. As shown in Table 2, jobs that respondents believed “did not fit well at all” took them, on average, more than twice as long to make the assignment than jobs respondents thought fit “extremely well” (64.95 seconds vs. 28.93 seconds). A similar relation was observed between confidence and mean reaction time. Moreover, as shown in Figure 3, the relations between the *average* values of these three variables for each job are also quite strong. This is consistent with the conclusion that these are all adequate measures of the goodness of fit of the jobs to the SOC. Furthermore, the correlations in Table 3 show that the composite score was very strongly positively correlated with its four constituents, reinforcing the conclusion that these were all convergent indicators of goodness of fit.

The final correlation at the bottom of Table 3 shows that the proportion of participants who said a job was new and emerging was correlated slightly negatively with the composite goodness of fit score ($r = -.21$). This means that new and emerging jobs did not fit the SOC as well as jobs that were not new and emerging. Participants manifested

Table 1. Correlation matrix of compatibility criteria

	Agreement	Fit	Confidence	Reaction time residuals
Agreement	1.00	–	–	–
Fit	.51	1.00	–	–
Confidence	.50	.86	1.00	–
Reaction time residuals	–.40	–.43	–.53	1.00

Table 2. Mean reaction times (sec.) for various assessments of fit and confidence

	Fit	Confidence
5	28.93 (81)	32.68 (99)
4	32.86 (361)	34.46 (554)
3	45.43 (423)	44.06 (504)
2	50.41 (463)	49.97 (243)
1	63.95 (220)	88.14 (148)

Note: Ns in parentheses. 5 indicates that the job “fits extremely well” or that the respondent is “extremely confident.” 1 indicates that the job does “not fit at all” or that the respondent is “not at all confident.”

considerable agreement about goodness of fit and confidence. At least four out of nine participants assigned the same goodness of fit for 129 jobs and the same level of confidence for 133 jobs.

To separate the 172 jobs into two groups, those that did and did not fit the SOC adequately, we identified the jobs whose composite scores were more than one standard deviation below the mean (< -3.23). This bifurcation produced 40 jobs (out of 172) that seemed not well represented in the SOC. We discarded the remaining 132 jobs as they seemed well represented and did not necessitate the formation of additional occupational categories (see the left-hand side of Figure 2).

We used data from the sorting task to separate these 40 jobs into groups. For each job, we recorded the labels of the piles in which the nine sorting task participants placed it. If a majority of participants agreed on the pile label for a job, we considered the job to be a member of the group identified by the pile label. Below, we refer to these descriptions as *consensual pile labels*. If two participants’ pile labels generally reflected the same type of occupation, we considered them to be in agreement. For example, we recorded “R & D” and “Scientist” as common pile labels since the two titles reflect the same underlying type of occupation: discovery research in the laboratory setting. Similarly, “Clinical Development,” “Clinical/Regulatory Affairs,” and “Regulatory” are considered to be in agreement because all of the three titles refer to occupations dealing with the clinical trials process and getting new discoveries approved by government regulators. It is important to point out that the interpretation of the pile labels is not algorithmic and requires some judgment on the part of the practitioner.

Table 3. Correlations between composite compatibility score and compatibility criteria

	Composite compatibility score
Agreement	.74
Fit	.87
Confidence	.89
Reaction time residuals	-.73
New and emerging	-.21

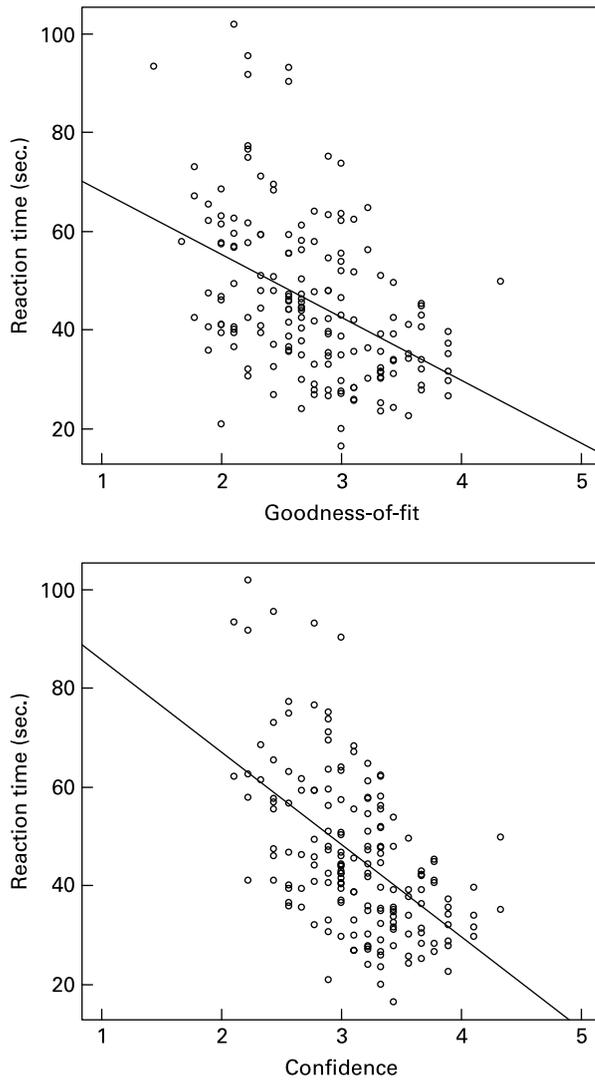


Fig. 3. Relationship between reaction time (sec.) and goodness-of-fit/confidence

For 21 of these jobs, consensus pile labels existed. For the remaining 19 jobs, our participants evidenced no consensus about the proper category to represent them (i.e., a majority of participants did not share a common pile label), and they were discarded (see middle of Figure 2). We used the pile label data from the sorting task to place these 21 jobs into nine groups. For five groups, multiple jobs consensually fit together (see right-hand side of Figure 2). In the first group, the individual pile labels identified by a majority of participants all contain the phrase “Clinical,” which we used as the consensus pile label. Similarly, each of the majority pile labels in the second group deals either with quality assurance or quality control (or both), suggesting that “QA/QC” was the most appropriate consensus pile label. The majority pile labels of the third group all contain the phrase

“Manufacturing,” which we used as well. In the fourth group, all of the highlighted labels contained the word “Regulatory,” which is the consensual pile label we selected as well. Finally, for the two jobs in the fifth group, the phrase “Project Manager” or “Project Management” is shared by all of the majority pile labels. We therefore chose “Project Manager” as the consensus pile label for this group of jobs.

In addition, four jobs each consensually fell under their own pile label (see right-hand side of Figure 2). For each of them, a majority of participants agreed upon the same phrase to write on the label. We used these same phrases to construct the following four consensual pile labels for jobs six through nine, respectively: “Business Development,” “Facilities,” “IT,” and “Customer Support.” The nine consensus pile labels of the incompatible occupational groups are presented in the top panel of the left-hand side of Table 4.

After identifying these general categories, we examined their associated Radford job titles and descriptions to construct possible additions to the SOC. For the “Clinical” group, the Radford descriptions depicted jobs dealing with the management and oversight of the clinical trials process. Thus, an appropriate addition to the “Management” major group of the SOC would be the detailed occupation: “Clinical Trials Managers.” For the jobs in the “QA/QC” category, the job descriptions all dealt with management over quality control and assurance functions. Thus, a potential SOC category that encompasses these jobs is “Quality Control Managers,” which belongs under the “Management” major group. For the

Table 4. Proposed SOC additions according to three analytical techniques

	Consensus pile label	Proposed SOC addition
<i>Incompatible occupations</i>		
1	Clinical	Clinical trials managers
2	QA/QC	Quality control managers
3	Manufacturing	Manufacturing process technicians
4	Regulatory	Quality control auditors
5	Project manager	Project managers
6	Business development	Business development analysts
7	Facilities	Facilities designers
8	IT	Biomedical systems analysts
9	Customer support	Product complaint representatives
<i>New and emerging occupations</i>		
1	Clinical affairs/regulatory	Clinical trials managers
2	Manufacturing/process development	Biomedical process and development engineers
3	Environmental health and safety/facilities	Toxicologists
4	Business development	Clinical research managers
5	IT/automation	Automation engineers
<i>Clustering analysis</i>		
1	Clinical development/operations	Clinical trials managers, analysts, and technicians
2	Manufacturing/process and product development	Manufacturing process designers and coordinators
3	Quality assurance/quality control	Quality control managers
4	Business development/contracts management	Business development managers

“Manufacturing” category, the Radford titles and descriptions all depicted occupations dealing with the design of manufacturing and production processes as well as the technical support associated with those processes. Therefore, “Manufacturing Process Technicians” could be added to the “Production” group in the SOC. Jobs in the fourth category (“Regulatory”) also dealt with quality control and assurance issues, but they had more to do with audit functions (as opposed to management). Thus, an additional QA/QC category can be added to the SOC, but in the “Production” major group: “Quality Control Auditors.” The two job titles in the “Project Manager” group were basically identical to the pile labels. As a result, the SOC category “Project Manager” could be added to the “Business and Financial Operations” group. This translation of the Radford job descriptions into occupational titles suitable for the SOC is another instance in which judgment is required by the practitioner.

The remaining four categories each only contained one job. We therefore used the Radford titles and descriptions to make recommendations for additional categories to be added to the SOC. The proposed SOC revisions associated with the nine consensual pile labels are presented in the top panel of Table 4.

5.2. Analysis 2: New and Emerging Occupations Not Represented in the SOC

Whereas the first method relied on implicit measures of compatibility, the second method determined which jobs in biotechnology were new and emerging using explicit reports. We again started with the entire set of 172 job descriptions, as shown in Figure 4, which illustrates the steps of the second method. A majority of participants (at least five out of nine) identified 31 jobs (out of 172) as “new and increasing” in response to the third question of the SOC assignment task. The remaining 141 jobs were discarded as they represented long-standing work done in the field. For 27 of these jobs, consensus pile labels existed (refer back to the definition of “consensus pile label” in Subsection 5.1). As with the previous analysis, we used the pile label data from the sorting task to place these 27 new and emerging biotechnology jobs into six groups. For the remaining four jobs, our participants evidenced no consensus about the proper category to represent them (i.e., a majority of participants did not share a common pile label), and they were discarded.

23 of the 31 jobs fit into one of the two broad categories described above: “Research & Development Scientist” and “Clinical Affairs/Regulatory.” Four other job descriptions

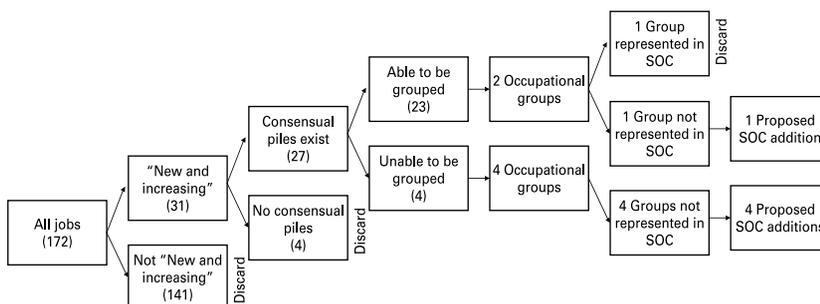


Fig. 4. Flowchart of identification of “New and Emerging” occupations

each fell into their own categories. The individual pile labels associated with one job reflected jobs dealing with the process of producing and manufacturing approved therapies. We therefore constructed the title: “Manufacturing/Process Development.” For the second job, almost all of the participants’ pile labels reflected occupations related to managing and maintaining facilities and making them compliant with occupational health and safety regulations. This led us to create the title: “Environmental Health & Safety/Facilities.” For the third job, eight of the nine participants all wrote down essentially the same pile label: “Business Development.” We use this title as well. Finally, for the fourth job, the majority of participants wrote pile labels reflecting information technology work related to designing automation processes. Consequently, the consensus pile label is: “IT/Automation.”

We next determined whether any of the new and emerging occupational categories fit into existing SOC categories. We assumed that participant agreement about assignment of a job to a particular SOC category would indicate that the job description adequately fits into the SOC. We therefore gauged whether a substantial plurality of participants (at least four out of nine) assigned each new and emerging job description to the same SOC code. This cut point was selected because the SOC categories chosen for job descriptions regarding which four or more respondents agreed had extremely high face validity with the Radford job titles. In other words, the wordings of the SOC category and the Radford title sounded very similar for jobs that exceeded this threshold. Those below this cut point were notably less valid. This selection criterion admittedly requires judgment and domain-specific knowledge on the part of the researcher, but the effective revision of classification systems necessitates both quantitative and qualitative analysis.

Participants were able to agree on the placement of many of the job titles in the “Research & Development Scientist” group, but there was a lack of consensus with respect to the titles in the other groups. Therefore, we concluded that “Research & Development Scientist” was already represented in the SOC, and we accordingly discarded the group as shown on the right-hand side of Figure 4. However, “Clinical Affairs/Regulatory,” “Manufacturing/Process Development,” “Facilities/Environmental Health and Safety,” “Business Development,” and “IT/Automation” were not represented in the SOC and represent categories of jobs that might be added (see the middle panel of the right-hand side of Table 4).

After we identified the five general categories of new and emerging jobs that were not compatible with the SOC, we examined the Radford job titles and descriptions associated with those categories to construct possible additions to the SOC. As mentioned earlier, the “Research & Development Scientist” group was well represented by the SOC and did not require a new SOC category. In the “Clinical Affairs/Regulatory” group (which was found to be incompatible with the SOC), the Radford descriptions depicted jobs dealing with the management and oversight of the clinical trials process. Thus, we concluded that an appropriate addition to the “Management” major group of the SOC would be the detailed occupation: “Clinical Trials Managers.”

The other four categories only contained one job; therefore, we could simply use the Radford job title and description as the proposed SOC addition. The five recommended additions to the SOC (and their associated consensus pile labels) are presented in the middle panel of Table 4.

5.3. Analysis 3: Identifying the Optimal Structure of Biotechnology Jobs

The final analysis we conducted took a different approach and attempted to identify the optimal organization of the jobs in biotechnology without use of the SOC. In other words, we sought to identify the best scheme for organizing these jobs into superordinate categories. Ideally, such a scheme would be the SOC (perhaps with modifications). The steps of the third method are illustrated in Figure 5.

For this analysis, we analyzed the data from the sorting task using cluster analysis to determine which jobs fell in the same “clusters” (i.e., organizational categories). For each of the nine participants, we produced a 172×172 matrix with the job descriptions as the rows and columns. In each cell, we entered a 1 if the participant put the two jobs in the same pile, a 0 if he or she did not, and 1's along the diagonal. Therefore, the matrix was symmetric. We then combined the individual matrices into a 172×172 symmetric distance matrix for the entire sample that contained information about the level of dissimilarity between each pair of jobs. In each cell, we entered $(1-X)$, with X being the proportion of participants (out of nine) who put the two jobs in the same pile. Hence, the diagonal entries of the distance matrix are 0, and the off-diagonal entries represented the amount of “distance” between each pair of jobs.

We then applied three different hierarchical clustering algorithms to the distance matrix – average linkage, weighted average linkage, and Ward's method.⁴ For each of the three methods, we applied the Calinski and Harabasz (1974) stopping rule to determine the correct location to cut the dendrogram.⁵ For all three approaches, the optimum number of clusters turned out to be 15, as shown on the left-hand side of Figure 5. In general, the three methods placed similar jobs in similar categories. When disparities occurred, we assigned a job to a cluster only if two out of the three approaches identified the job as belonging to the cluster. This procedure for handling discrepancies had no substantive impact on the findings, because only a small minority of jobs in each cluster was anomalous.

As illustrated in the middle of Figure 5, we coded two sets of information for each cluster: cluster labels and the SOC codes associated with each cluster. First, the information from the pile labels was used to label each cluster. This was done by first looking up the pile label associated with each job in a given cluster for each of the nine participants. If a majority of participants placed a majority of the jobs within the same cluster, then we used their descriptions for the cluster label. This procedure produced at least one cluster label for every cluster. Second, we used information from the assignment task to determine the SOC codes associated with each cluster. We assigned a SOC code to a cluster if for any of the jobs within that cluster, four or more respondents expressed agreement.

⁴ Clustering involves reducing data into smaller groups. In order to this, one must define the distance between a point and a cluster and the distance between clusters. There are several procedures for doing this. The average linkage method defines distance as the average distance between all pairs of points; the weighted average linkage method is similar, except that it weights each *group* of observations (as opposed to each observation) equally. Ward's method groups observations to minimize information loss (as measured by the total sum of squared deviations of every observation from the mean of its cluster) in each partition of the data. For more information about clustering techniques, see Dillon and Goldstein (1984).

⁵ The Calinski and Harabasz (1974) stopping rule employs pseudo-*F*-tests to determine whether an additional partition significantly reduces within-cluster dissimilarity. It is generally considered to be the best-performing across a wide range of applications (Milligan and Cooper 1985; Tibshirani et al. 2001).

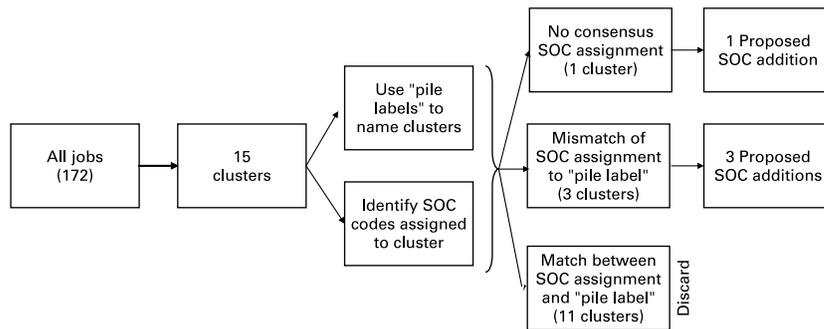


Fig. 5. Flowchart of cluster analysis identifying the optimal structure of biotechnology jobs

The adequacy of the SOC could be assessed by comparing the labeling of the clusters from the sorting task to the SOC assignments of those same clusters. Two circumstances raised “red flags.” First, if no consensus appeared regarding SOC code(s) for a cluster of jobs, then that would suggest that the SOC does not accurately reflect that occupational category. As shown in Figure 5, one labeled cluster did not have an SOC code associated with it (“Manufacturing/Process and Product Development”). Second, even if there was consensus regarding the appropriate SOC code, the SOC titles may differ considerably from the cluster labels given by our sorting participants. Three clusters had SOC codes that did not match the cluster labels (the second “red flag”): “Clinical Development/Operations,” “Quality Assurance/Quality Control,” and “Business Development/Contracts Management.” In these cases, there may be agreement about the closest SOC code, but that code may still not accurately reflect the specifics of the set of job descriptions that were clustered. Using the general cluster labels identified as “red flags” as well as their associated Radford titles and descriptions, we constructed four possible additions to the SOC, which are presented in the bottom panel of Table 4. As illustrated in the right-hand side of Figure 5, we discarded the remaining eleven clusters for which there was a match between the cluster label and the assigned SOC codes.

5.4. Summarization

Finally, we summarized the findings from the three analyses to identify the principal omissions in the classification system, as shown on the right-hand side of Figure 1. It is reassuring to note the convergence of findings across the three independent efforts at identifying categories that might be added to the SOC to improve its accommodation of jobs in biotechnology. One way to do so is shown in Table 5, which lists the detailed job titles that could be added to the SOC as indicated by the full analysis of data from the SOC assignment task (shown in Column 1), the “new and emerging” approach we took (shown in Column 2), and the full analysis of the sorting data (shown in Column 3). Given that the three analyses tackle the question from different angles, it is reasonable to assume that if a proposed SOC addition is implicated by a majority of the analyses, then that recommendation warrants special attention.

The top four panels of Table 5 display categories that might merit addition as “Detailed Occupations” according to two or three of the analyses. Four job titles were found not to be represented in the current SOC according to this approach: “Clinical Trials Managers,”

Table 5. Summary of results (Recommended additions of detailed occupations to the SOC as suggested by three analytic techniques)

Incompatible occupations	New and emerging occupations	Cluster analysis
Clinical trials managers	Clinical trials managers Clinical research managers	Clinical trials managers, analysts, and technicians
Manufacturing process technicians	Biomedical process and development engineers	Manufacturing process designers and coordinators
Quality control managers Quality control auditors		Quality control managers
Business development analysts		Business development managers
Biomedical project managers Facilities designers Biomedical systems analysts Product complaint representatives	Toxicologists Automation engineers	

Note: Table lists “Detailed Occupations.” The “Major Groups” in which these would be placed are discussed in the text.

“Manufacturing Process Designers and Technicians,” “Quality Control Managers,” and “Business Development Managers and Analysts.” The bottom panel of Table 5 displays the detailed occupations that emerged in only one of the three sets of analyses. Due to the lack of convergence across assessment methods, perhaps it is wisest not to consider them for addition to the SOC.

It may appear as if some of these proposed additions are already represented by the current SOC. For example, some may argue “Quality Control Managers” fit under the detailed occupations “Production, Planning, and Expediting Clerks” or “Weighers, Measurers, Checkers, and Samplers, Recordkeeping.” However, our participants disagreed with one another in terms of which specific SOC category to which to assign the biotechnology jobs dealing with QA/QC and believed that these jobs were incompatible with the existing SOC. This indicates that QA/QC is done in biotechnology uniquely enough to warrant representation in a separate SOC category.

Although the three approaches generally converged, there were meaningful differences between them. As seen in Table 5, the “new and emerging” occupations approach appears to have been more exclusive with respect to the number of revisions to make to the SOC, whereas the “incompatible occupations” approach appears to have been more inclusive. On the other hand, the clustering analysis did not “miss” any of the four occupations described above, nor did it identify a large number of occupations unique to itself. Hence, if practitioners were forced to choose only one analytic technique, our case study suggests that the clustering analysis might be optimal.

6. Recommendations for Researchers

We have described two cognitive exercises and three analytic techniques for identifying omitted categories in classification systems. Our case study applying this approach identified four occupations that could be added to the SOC to more adequately represent jobs in biotechnology. Considering that the SOC currently contains 822 detailed occupations, adding four for just one field of work is not trivial. But the fact that this is a relatively small number attests to the effectiveness of the SOC in its current form. Adding these four new categories would slightly decrease the counts of workers in other occupations in which many biotechnology jobs have been misclassified. Our analyses suggest that up to 40 unique biotechnology jobs (representing more than 23% of the total jobs) have not been well represented by the SOC and would be better represented by the four recommended additions.

Our revised version of the SOC would be more consistent with how biotechnology experts perceive the occupational structure of their field and would more effectively help the industry in its hiring practices and relations with government. However, any evaluation of the efficacy of the revised SOC in applications (such as the formation of education or immigration policies) would need to be assessed by other criteria. For instance, research could evaluate whether revising the SOC improves inter-rater agreement by government coders of biotechnology jobs or helps to optimize the number of H1-B visas that are approved to augment the workforce as needed.

The procedures we proposed for revising the SOC can, in principle, be exported to other industries represented in the SOC and to classification systems other than the SOC, but there are practical challenges involved. First, a suitable set of stimuli must be identified to

provide to respondents. In nearly every major industry, consulting firms or business associations exist that have compiled lists of occupations, although perhaps not as detailed and comprehensive as Radford's. Outside the occupational domain, researchers can leverage responses to open-ended questions in representative sample surveys of businesses and individuals. When no off-the-shelf stimuli are available, additional effort must be invested in conducting surveys or extensive focus groups of experts. These tasks may also need to be undertaken when existing materials lack necessary detail.

Second, many practitioners may not have access to the laboratory space and software required to administer the assignment task. Given the convergence of findings across the methods we examined, researchers could choose to perform only the sorting task, which does not necessitate any specialized equipment. In our study, the sorting task was more cost-effective than the assignment task, because the former required four fewer participants and less time to complete. Considering that we paid experts \$75 per hour for their time, the differences in administration time and sample size requirements are not negligible. Also, the clustering analysis was the one approach that identified the four shared missing occupations but no others. Nevertheless, although the sorting task may be optimal for resource-constrained practitioners who can only feasibly implement one procedure, we encourage the use of multiple methods for the purposes of validation.

Third, a practitioner must have some knowledge of the domain being studied. Some aspects of the methodology are algorithmic, but others require more judgment. For instance, the labeling of the clusters, the construction of SOC titles, and the choice of thresholds based on face validity were all qualitative decisions made by the researchers based on an examination of the data along with knowledge of the industry.

7. Coda

Governments and businesses depend on classification systems to produce statistics, distribute funds, implement policies, and assess performance. However, for these systems to be effective, they must be updated comprehensively and efficiently as the work environment changes. In this article, we have drawn on principles and techniques from cognitive psychology to offer one potential method of revision. The use of many different approaches to such a task is always desirable, and we hope this article has enriched practitioners' tool bags by adding psychological methods that are practical to implement and yield actionable results.

Appendix 1: Participant Recruitment

We sought human resources experts at biotechnology companies who could easily come to our laboratory. To identify such people, we created a list of major biotechnology firms in the San Francisco Bay Area using the member directory of the Biotechnology Industry Organization (BIO), the industry's principal trade association. BIO facilitates partnerships and joint ventures between member firms and lobbies on behalf of the industry. From the BIO website, we obtained the contact information for firms located in four telephone area codes (650, 415, 510, and 408), all within a 50-mile radius around Stanford University (which we thought was the maximum

distance participants would be willing to travel). Almost all biotechnology companies in the San Francisco and San Jose metropolitan statistical areas are located within those four area codes. Companies with fewer than 25 employees were eliminated, because human resources administrators at small companies might have limited expertise regarding the wide array of jobs in the field.

Each targeted company was telephoned and asked for the name and contact information of the person “who knows the MOST about the types of jobs being done by the people working at [Company X].” We did not specifically ask for human resources personnel, for two main reasons. First, smaller companies (25–50 employees) often have not established formal human resources departments, and many of the relevant responsibilities are assigned to professionals in departments of business development and administration and, in some cases, to CEOs. Second, in larger companies, many human resources managers serve administrative functions and do not have expert knowledge of the specifics of every occupation, particularly those in scientific fields.

Some companies refused to identify the relevant individual, but the majority of companies gave us names and phone numbers. We then contacted these individuals by telephone, described the details of the study we were conducting, and invited them to participate. Among the people we were able to contact, nearly all agreed to participate. This approach yielded 18 participants. Two additional participants were referred to us by Radford Surveys, and two more people were referred to us by an expert who had participated early on in the project.

This sample reflected important diversity in the field. The practice of biotechnology can be divided up into three major work phases (discovery, development, and commercialization), and our sample included individuals with expertise in all three of these areas. During the discovery phase, much of the research and development of therapies takes place. During the development phase, clinical trials take place in order to demonstrate the efficacy and safety of therapies to federal regulators. Commercialization involves the production and marketing of approved therapies. Of the nine firms represented in our sample for the sorting task, four engaged in discovery, eight engaged in development, and three engaged in commercialization. Commercialization firms are less common in the Bay Area than are discovery and development firms because of the high labor and facilities costs there. Of the 13 firms represented in the sample for the SOC assignment task, eight engaged in discovery, 13 engaged in development, and six engaged in commercialization.

Appendix 2: Instructions for the SOC Assignment Task

Thank you very much for your help with this research.

Your task will be to read 172 job descriptions and find the code that matches the description most closely.

The complete list of codes is on the large sheets that the researcher has given you.

When you read the job description, please be sure to consider all of the codes on the sheets, so you pick the one that matches the description most closely.

After you pick the best code, you will be asked how confident you are that you picked the best code and how closely the code matches the description.

In some cases, you may feel that a job description does not fit any of the categories. However, even in these cases, we need you to select the closest fitting category, even if it is a very poor fit, and indicate that in your rating.

Please work as quickly as you can without making any errors.

Please don't take breaks before the computer tells you to.

Before you begin, please take some time to familiarize yourself with the codes on the sheets.

Click the button below when you are ready to start.

Thank you!

Appendix 3: Instructions for the Sorting Task

The purpose of this session is to help us learn about how you think about jobs in the field of biotechnology, to help us understand how work is being done in America today and to help people outside your industry understand it better.

Everything you do today will be completely anonymous. Please don't write your name on any materials I give you.

What I'd like to ask you to do is what we call a "sorting task." I'm going to give you a set of cards. On each one is a description of a job in a biotechnology company. The description often lists the tasks an individual is expected to perform in his or her work.

Please read through the stack of cards and sort them into piles. Put job descriptions that seem to go together in the same stack, and create a different stack for each different category of jobs. You can create as many or as few piles as you think are needed.

We want to learn how you naturally organize the jobs and see similarities and differences among them.

When you finish sorting the piles, it would help me if you would please put one of these blank cards on top of each pile and write on each one a label for the pile to identify what type of jobs you put in that pile. Please do not write anything on the cards themselves. You can fill out these label cards at any time, just as long as every pile has a label on top when you're done sorting.

While you're working on this, I will be next door. You're welcome to check with me any time if I can do anything to help you.

Please take your time to do this carefully so that the end result you produce accurately represents how you think about these jobs. And as you go along, if you decide you want to change your piling system, feel free to do so as much as you like.

Do you have any questions before you start?

Appendix 4: Reaction Time Regressions

Calculation of Corrected Reaction Times

In this section, we explain how we calculated the adjusted reaction time scores. The computer program we developed recorded to the microsecond how long it took a participant to select an SOC code to match a given job description. These times are likely to be a function not only of the ease of finding an SOC code for a job but also of three other factors as well: length of job description, practice, and baseline speed differences between

people. To remove the effect of these latter three factors from our reaction time scores, we conducted a multiple regression and calculated residual scores.

The parameters of the following linear regression equation were estimated using OLS to predict reaction time:

$$T_i = \beta_0 + \beta_1 P_i + \beta_2 P_i^2 + \beta_3 L_i + \beta_4 L_i^2 + \lambda \sum_{r=1}^{12} R_{ri} + \varepsilon_i \quad (1)$$

where T_i represents the natural log of the reaction time in seconds for observation i , P_i represents the position in the computer exercise (1–172) of the job description as seen by the participant, L_i represents the length of the job description in sentences, R_{ri} are dummy variables for participants $r = 1, \dots, 12$, and ε_i represents random error. Because reaction time data are often skewed, we normalized the values of the dependent variable by taking their natural logs (Fazio 1990). We also included squared terms for position and length to allow for nonlinearity. We assumed that the construct we are trying to measure is captured by the error term. Hence, we obtained corrected reaction time scores by calculating the average residuals of the regression ($T_i - \hat{T}_i$) for each job description across all thirteen participants. In order to maintain consistency with the other statistics, we used negative residuals so that higher values indicate quicker reaction times.

Validity of the Reaction Time Measurement

We conducted an additional regression analysis to gauge the validity of the reaction time measurements:

$$T_i = \beta_0 + \beta_1 P_i + \beta_2 P_i^2 + \beta_3 L_i + \beta_4 L_i^2 + \beta_5 F_i + \beta_6 C_i + \lambda \sum_{r=1}^{12} R_{ri} + \varepsilon_i \quad (2)$$

This is the same equation as above, with two additional predictors included: goodness-of-fit ratings (F_i) and confidence ratings (C_i). The parameter estimates for this equation appear in Column 1 of Table 6. As expected, we found an inverse relation between position and reaction time, meaning that participants were able to more quickly match job descriptions presented later on in the exercise. Nonlinearity in this relation indicated that the greatest practice effects occurred early on in the sequence of ratings. We also saw a positive relation between job description length and reaction time; longer job descriptions yielded longer reaction times. This relation was also nonlinear, with larger increases in time coming from increases in paragraph length nearer to zero.

Of the twelve dummy variables representing the 13 participants, three of these were nonsignificant, and nine were significant, meaning that there were reliable differences between participants in their baseline speed of responding.

Confidence ratings had a significant, negative relation with reaction time. This suggests that people took longer to assign codes when there were multiple contenders among the SOC codes to represent a given job.

As expected, the coefficient for goodness-of-fit ratings was negative, meaning that people took longer to assign a SOC code when they felt that the code did not fit the job description well. Although this regression coefficient was not statistically significant, this may have

Table 6. Regressions predicting natural log of reaction time (sec)

	1	2
Paragraph position	-.02***	-.02***
Paragraph position squared	6.4×10^{-5} ***	6.3×10^{-5} ***
Paragraph length (sent.)	.09*	.06*
Paragraph length squared	-.01**	-.01**
Goodness-of-fit	-.02	-.19***
Confidence	-.24***	-
Participant 1	-.37**	-.07
Participant 2	-.89***	-.49***
Participant 3	-.54***	-.01
Participant 4	-.22*	-.04
Participant 5	.07	.28*
Participant 6	-.78***	-.66***
Participant 7	-.74***	-.45***
Participant 8	-.67***	-.39**
Participant 9	-.18	-.22 ⁺
Participant 10	-.70***	-.39**
Participant 11	.09	.34**
Participant 12	-.62***	-.33**
Constant	5.36***	4.84***
R ²	.514	.474
Adj. R ²	.509	.469
N	1,548	1,548

*** $p < .001$; ** $p < .01$; * $p < .05$ (two-tailed)

been the consequence of multicollinearity, as fit and confidence are highly correlated with one another. When removing confidence from the specification, goodness-of-fit emerges as a negative and highly significant predictor of reaction time (see Column 2 of Table 6).

In general, this regression suggests that the reaction time measurement behaved as expected, which reinforces our confidence in its validity.

8. References

- Brewer, M.B. and Lui, L.N. (1989a). Use of Sorting Tasks to Assess Cognitive Structure. In *Answering Questions: Methodology for Determining Cognitive and Communicative Processes in Survey Research*, N. Schwarz and S. Sudman (eds), San Francisco: Jossey-Bass Publishers.
- Brewer, M.B. and Lui, L.N. (1989b). The Primacy of Age in the Structure of Person Categories. *Social Cognition*, 7, 262–274.
- Calinski, T. and Harabasz, J. (1974). A Dendrite Method for Cluster Analysis. *Communications in Statistics*, 3, 1–27.
- Chi, M.T., Feltovich, P.J., and Glaser, R. (1981). Categorization and Representation of Physics Problems by Experts and Novices. *Cognitive Science*, 5, 121–152.
- Coxon, A.P.M. (1999). *Sorting Data: Collection and Analysis*. In Sage University Paper Series on Quantitative Applications in the Social Sciences, 07–127. Newbury Park, CA: Sage.

- Dillon, M. and Goldstein, W.R. (1984). *Multivariate Analysis: Methods and Applications*. New York: John Wiley and Sons.
- Donders, F.C. (1969). On the Speed of Mental Processes. *Acta Psychologica*, 30, 412–431.
- Fazio, R.H. (1990). A Practical Guide to the Use of Response Latency in Social Psychological Research. In *Review of Personality and Social Psychology*, C. Hendrick and M.S. Clark (eds). Newbury Park, CA: Sage.
- Kelleher, M., Scott, P. and Jones, B. (1993). Resistant to Change? Some Unexplained Omissions in the 1990 Standard Occupational Classification. *Work, Employment and Society*, 7, 437–449.
- Levine, C., Simon, L., and Weinberg, D.H. (1999). Revising the Standard Occupational Classification System. *Monthly Labor Review*, 122, 36–42.
- Milligan, G.W. and Cooper M.C. (1985). An Examination of Procedures for Determining the Number of Clusters in a Data Set. *Psychometrika*, 50, 159–179.
- Pollack, L.J., Simons, C., Romero, H., and Hausser, D. (2002). A Common Language for Classifying and Describing Occupations: The Development, Structure, and Application of the Standard Occupational Classification. *Human Resource Management*, 41, 297–307.
- Rips, L.J. Shoben, E.J., and Smith, E.E. (1973). Semantic Distance and the Verification of Semantic Relations, *Journal of Verbal Learning and Verbal Behavior*, 14, 665–681.
- Rosch, E. (1975a). Cognitive Reference Points. *Cognitive Psychology*, 7, 532–547.
- Rosch, E. (1975b). Cognitive Representations of Semantic Categories. *Journal of Experimental Psychology: General*, 104, 192–233.
- Rosch, E., Simpson, C., and Miller, R.S. (1976). Structural Bases of Typicality Effects. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 491–502.
- Salmon, L. (1999). 1999 Standard Occupational Classification. *Industrial Relations*, 38, 231–235.
- Schacter, D.L (1987). Implicit Memory: History and Current Status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 501–518.
- Schuman, H. and Presser, S. (1996). *Questions and Answers in Attitude Surveys*. Thousand Oaks, CA: Sage.
- Tibshirani, R., Walther, G., and Hastie, T. (2001). Estimating the Number of Clusters in a Dataset Via the Gap Statistic. *Journal of the Royal Statistical Society, Series B*, 63, 411–423.
- Welford, A.T. (1980). Choice Reaction Time: Basic Concepts. In *Reaction Times*, A.T. Welford (ed), New York: Academic Press.
- U.S. Bureau of Labor Statistics (2004). *Standard Occupational Classification (SOC) User Guide*. Washington, DC: Department of Labor.

Received October 2005

Revised March 2007