Deskilling revisited: new evidence on the skill trajectory of the Australian economy 2001-2007

Doug Fraser

Visiting Research Fellow, Australian Defence Force Academy, University of New South Wales
Research Associate, Industrial Relations Research Centre, University of New South Wales
Research Associate, Australian Innovation Research Centre, University of Tasmania

Paper delivered to the International Labour Process Conference

Rutgers University, New Brunswick NJ

15 March 2010

ABSTRACT

This paper describes a new common metric capable of comparing changes over time in the skill content of jobs across the full spectrum of industries, occupations and occupational levels. Based loosely on Spenner's model and the indices used in the ESRC Skills Surveys, this metric combines two constructs: skill-intensity (the degree to which a job stretches the skill base of the worker) and task discretion (the amount of choice which workers, individually or collectively, exercise over what they do in their job, how and when). The metric is applied using data from the first seven waves of HILDA, a multi-purpose annual panel survey with a large nationally representative sample, to track the skill trajectory of the Australian economy.

The data up to this point suggest that on average, the overall skill content of Australian jobs fell over the first five years of the 21st century, with signs of renewed growth from 2005 onwards. The pattern for task discretion is mixed, with overall declines in employees' control over the tasks done and in participation in decision-making, but a steady though small improvement in control over the timing of work. Moreover, some important industries show alarming discrepancies between the skill content of work and the control workers have over their jobs. Factor analysis of the response suggests that when Australian workers assess the quality of their jobs, task discretion - in particular, control over the timing of their work - is a far more salient issue than the need or opportunity to exercise skill. Together, these findings suggest a need to reconsider Spenner's "pragmatic hypothesis".

Deskilling revisited: new evidence on the skill trajectory of the Australian economy 2001-2007

Doug Fraser[1]

Introduction

The original deskilling debate of the 1970s and 80s ran out of steam somewhere around 1990, largely because there were no means at that time of demonstrating conclusively whether the kinds of experience that had been so persuasively documented over a large body of case studies were part of a secular or dominant trend (Attewell 1987; Form 1987). The type of quantitative research which would have been needed to convince anyone, either way, was ultimately ruled out by the absence of any large-scale data source addressing skill, except through the crudest of proxies; these in turn generally made it impossible to distinguish skill from confounding factors, such as social prestige and educational prerequisites, which often begged the very questions the research was intended to answer.

This data gap was presumably responsible for the development of purpose-designed sample surveys aimed at tracking skill through common but specific indicators, of which the most comprehensive so far have been the series of five loosely coordinated surveys conducted in the UK since 1986 (Felstead, Gallie and Green 2002, 2004; Felstead, Gallie, Green and Zhou 2007). While this research has proved extremely informative about the skill trajectory of the UK economy, it has been difficult so far to extend it to other nations because they lack data sources of similar quality and comprehensiveness. This paper seeks to further such research by proposing a new basis for tracking changes over time in the skill content of jobs[2], whether individually or across an industry or a national economy, which is significantly less data-intensive but preserves the same advantages of neutrality and specificity.

The metric I describe was designed to fulfil four key needs:

i. to measure skill *as actually exercised*, rather than the amount of *potential* productivity present in the economy in the form of workers' (or potential workers') qualifications, capabilities and experience;
ii. to capture the element of dynamism in the skilling system (i.e. flows as opposed to stocks) and hence be sensitive to small variations over short intervals of time;
iii. to capture as much as possible of the variety which actually exists, rather than suppress variation by averaging out individual observations or reducing them to job categories;
iv. to be uniformly applicable across the full range of industries and occupational hierarchy, without bias.

These four criteria differ from those of many metrics applied in earlier studies of national skill trajectories (Spenner 1983, 1985, 1990; Attewell 1990; Felstead, Gallie and Green 2002), and reflect the origins of my current research in innovation studies. My underlying purpose has been to assess the

---

[1] Visiting Research Fellow, Australian Defence Force Academy, University of New South Wales. The original research on which this paper is based was carried out under the auspices of the Australian Innovation Research Centre, University of Tasmania.

[2] In this paper the word "job" is used in a micro sense to describe the specific work done by an individual worker in a given role, in a given workplace, at a given point in time. This definition needs to be distinguished from the more common usage which averages out the experiences of individual workers into standardised *categories* of job, e.g. the duties conventionally assigned to a particular occupational classification in a particular industry or firm.

contribution of skill to an industry's or a nation's potential to innovate successfully - skill as a factor in economic dynamism and adaptiveness. The model of innovation I follow is not one based on big bangs or high-end breakthrough technologies, but an incremental approach which sees even major transformations as starting small and local, spreading unevenly and idiosyncratically to the point where they become pervasive, and benefiting from continuous micro-scale adaptation, learning and re-invention at each point in the diffusion process (Fagerberg 2004; Hall 2004). Such a distributed model of innovation implies a skill requirement going beyond just a core of leading-edge technological skills, with the implementation of change depending on the presence of appropriate knowledge, capabilities and aptitudes (not to mention incentives) at many points in the occupational structure: hence the need to capture even small movements, over small intervals of time, which might presently develop into cascades, and to map the growth and decay of skills indiscriminately across the workforce, without waiting for it to become apparent in broader aggregates such as standard job descriptions, and without unduly privileging the sexy or highly knowledge-intensive areas of work.

While this approach necessarily leads to different emphases from those of more traditional approaches to deskilling based on class analysis or industrial relations, I would nonetheless argue that it ends up on common ground. In the first place, following the now accepted practice in innovation theory (Edquist 2004, Malerba 2004), this analysis takes place in a system framework where outcomes are seen as arising out of multiple two-way interactions between a wide variety of infrastructural, behavioural and institutional factors, rather than assuming a simple linear demand-supply-output chain. Industrial relations (or employment law more generally) are relevant to this model precisely because they are one of the more important systems influencing the National System of Innovation. But more importantly, this approach implies a view of the productive exercise of skill as being the combined outcome of multiple interactions between supply, demand and deployment. That places my research squarely in that important tradition of labour process theory which derives from the skill equilibrium (Finegold and Soskice 1988, Wilson and Hogarth 2003) and skill ecosystem (Finegold 1990, Buchanan 2006, Hall and Lansbury 2006) models.

I must also confess that my choice was partly a matter of the data I could lay my hands on. While the metric I set out here lends itself ideally to application with large-scale microdata, few independent researchers have the clout or resources to generate such data to order; we must rely, most of the time, on found objects. My own data source, the Household, Income and Labour Dynamics in Australia (HILDA)[3] survey, was fortuitous in just that way, and I had little option but to make the most sense I could out of the indicators that were there. The empirical limitations of my metric owe as much to the restricted set of skill-relevant variables, and to the peripheral status of those variables within the dataset as a whole, as to any deliberate methodological choice.

---

[3] This paper uses unit record data from the Household, Income and Labour Dynamics in Australia (HILDA) survey. The HILDA project was initiated and is funded by the Australian Government Department of Families, Housing, Community Services and Indigenous Affairs (FaHCSIA) and is managed by the Melbourne Institute of Applied Economic and Social Research (MIAESR). The findings and views reported in this paper are, however, those of the author and should not be attributed to either FaHCSIA or the MIAESR.

Measuring skill

One reason the original debate never came close to empirical resolution was that there was so little agreement between authors on what they meant by skill (Attewell 1990). To oversimplify grossly, most of the earlier works that sought to provide evidence of deskilling, in a tradition dating back to Marx and even to Adam Smith, concentrated primarily on the loss of worker autonomy. Most of the work by economists, generally setting out to demonstrate the opposing case, sought to track the objective skill content of jobs, whether directly through sources such as the Dictionary of Occupational Titles, or through indirect indicators relating to such things as pay classification, years of education or required qualifications. It is hardly surprising that the two sides regularly carried on a dialogue of the deaf.

Spenner (1983, 1985, 1990) made repeated attempts through the 1980s to find some compromise position which would permit a debate on common evidentiary ground. Reviewing an already considerable body of literature, he suggested that both sides had a case, and hence that there were two dimensions of skill that should be regarded as "fundamental and underlying" in any operational definition of skill: *substantive complexity* (the level, scope and integration of mental, manipulative and interpersonal tasks in a job) and *autonomy/control* (the discretion or leeway available in a job to control the content, manner and speed with which tasks are done) (1985: 135; 1990: 402-3). While putting this simplified classification forward only as a "hypothesis and a pragmatic approach" (1990: 402), and not suggesting that this pair of dimensions captured all the significant variation, Spenner noted that it had the advantage of being "two primary dimensions of skill that have applicability across all jobs in the economy" (1985: 136).

The exact relation between the two dimensions has always been controversial. Spenner described them as "conceptually distinct but empirically correlated" (1985: 135), and later asserted that empirical studies in the literature had produced correlation coefficients in the range of .5 to .7 (1990: 403), a claim for which I have never been able to identify the original evidence and which has not been backed up by any subsequent research I know of, including my own. This became the main objection to his model, both at the time (Form 1987) and since (Lowry, Molloy and McGlennon 2008). Spenner's reaction to Form suggests a certain ambiguity in his own view. While taking care to distinguish autonomy/control (within a role) from authority (a relationship between roles), he ultimately concedes that it is immaterial whether autonomy/control is defined as part of skill, so long as it is treated as a relevant issue (1990: 405). In any case, the basic model has now amply demonstrated its practical value through its adoption by the UK surveys, and in the absence of any real advance on his suggestion since 1990, I have used it as the underlying framework for my own, with two very important changes.

Before proceeding to outline how my approach differs, it is worth reflecting a little on the two dimensions as Spenner saw them. Starting with substantive complexity, the three components of his definition can be read as having the following implications (though he himself never appears to have actually defined them). *Level* refers to the element of difficulty involved in carrying out a task, potentially including such things as knowledge content, required aptitudes, need for prior experience, responsibility and unpredictability. *Scope* refers not just to the number of different tasks comprehended by a job, but to the range of qualitative variation among them. A job that involves carrying out a great many separate tasks of the same kind (e.g. using a highly automated machine on a speeded-up assembly line, or handling heavy traffic in a call centre where the inquiries are routine and the interactions rigidly scripted) may be intensive or demanding simply because of the volume of work to be done in the time, but does not necessarily pass the test of complexity by this definition. *Integration* is the need and capacity to carry out tasks of different kinds, each requiring its own

distinctive knowledge and having its own distinctive success criteria, in a coordinated fashion so that they work synergistically towards a given end.

Complexity thus implies not just doing many things, or doing many things of different kinds, but doing them together to produce an outcome or output which is broader or sometimes different in scope from any one of them. The thing that makes it *substantive* is that those elements can be established objectively, by observation or measurement, rather than resting on the subjective assessment of the individual informant. This is important when considering the type and quality of evidence that is required to capture this dimension.

Autonomy/control is likewise a multidimensional construct. Spenner's definition has three elements: *content*, *manner* and *speed*. The first suggests that the worker has some choice in the selection and definition of the tasks which go together to make up the job. The second implies that even where the tasks themselves are predetermined, the worker has discretion to perform them in ways that are sensitive to the immediate situation or appropriate to her capabilities or knowledge. The third relates to control over the timing of the tasks. A satisfactory metric needs to capture each of these elements independently.

I retain this dimension in my model as *task discretion*, the same term used in the UK Skills Surveys, but with one very important difference. The British authors apply the term specifically to the control exercised by an individual worker over her own work, and are careful to distinguish this aspect of control from collective decision-making or consultation (Gallie, Felstead and Green 2004). I extend it to include any mechanism which enables the worker to feel that she has some control over the decisions affecting her immediate work, whether that mechanism consists of autonomous decision-making at the level of the individual, collective decision-making at the team or work-group level, or input to decision-making through some kind of formal or informal consultative process. In an increasingly interconnected world and workplace, I would argue that it is becoming harder to find jobs that involve a lot of genuine autonomy, and the most important issue today may be the ability to exercise some kind of collective control in a work setting characterised by high levels of interdependence. More generally, and given that this metric is based on worker self-report, I would argue that what really matters is the *sense* of being in control rather than the mechanism by which it is brought about.

I include task discretion largely because of its extensive logical overlap with substantive complexity. Increasing job complexity creates more situations where there is a range of choices for action, the consequences of a wrong choice are more critical and harder to reverse, and so many factors influence the correct choice that only someone on the spot has enough knowledge of the circumstantial factors to make a correct decision. The more such occasions can be expected to arise, the harder it becomes to codify the correct choices beforehand in a protocol or to reach them through micromanagement from a distance; consequently such jobs cannot be exercised effectively unless the jobholders, individually or collectively, have a high level of input to the relevant decisions.

Task discretion also interacts with the integration element of complexity. The skilfulness of a job, by this definition, depends to a large extent on whether the integration function is internal or external to the job. Where the worker is responsible for coordinating multiple tasks to produce a higher-order output, this in itself makes the job more complex and more skilful, and also implies greater control. But where the worker lacks this decision-making latitude, it is more likely that coordination will be handled externally: the worker will be expected to carry out a pre-specified set of tasks in a prescribed manner at prescribed times without regard to how they interact, and someone further up the line will take on the responsibility for integrating the outputs into a higher-level outcome.

Thus task discretion, even if not itself a characteristic that makes a job high-skilled, can be taken as evidence of its complexity, since highly complex jobs would not be sustainable unless they incorporated high levels of task discretion. In this sense I argue that it can represent a more or less acceptable *proxy* for substantive complexity in situations where there is insufficient evidence to track the latter.

This is precisely the problem that existed with my data. At least in the early waves, the only question in HILDA directly relating to complexity is one which asks the respondents for their *subjective* assessment of the complexity of their job. As noted above, this is a different matter altogether from *substantive* complexity, and cannot constitute reliable evidence of the latter in its own right. Later waves do contain questions on task variety, repetitiveness and the need to show initiative, which could in theory form the basis for a substantive complexity scale of sorts, but factor analysis of the response suggests that the respondents do not perceive them as part of a common construct.

Instead I have used as my complementary dimension one which is overlooked in Spenner's model, and which I call *skill-intensity*. Skill-intensity is the degree to which a job "stretches" the skill base of those who exercise it, independently of whether that skill base is high or low in its own right. The elements of "stretch" include the extent to which the worker uses all his existing skills in his work, the gap between his existing skills and the demands of the job, and the extent to which he wants, and is able, to learn new skills or refine his existing ones in order to do it properly or better.

Thus skill-intensity, like the other two dimensions, is multidimensional. Its three complementary sub-dimensions are *match*, *stretch* and *learning*. It is the combination rather than their individual presence that makes a job skill-intensive. A good match between the worker's skill base and the demands of the job, without the dynamic element of stretch, can lead to complacency, path-dependence and ultimately competence traps. Stretch is productive only if the workforce has the capacity and opportunity for learning to close up the skill gaps as they emerge. Learning, understood as both the perceived need to learn and the resources to fulfil that need, is worthwhile only to the extent that the job takes productive advantage of it within a reasonable time. Thus, an adequate indicator of skill-intensity needs once again to be a composite one, made up from multiple variables that capture the different aspects.

The particular strength of this construct lies in capturing the *dynamism* in the skilling system. It treats the skill content of a job, or the skill base of an individual, not as a fixed quantity existing in its own right but as something that grows, decays and changes constantly in line with the demands of the job, the informal learning that takes place on the job, the motivation of the worker and the management philosophy and work practices that prevail in the workplace. The mechanism I posit here is analogous to the one I suggested earlier as underlying innovation: change in the skill content of work does not (generally) occur because someone in authority declares unilaterally that the job description or the qualifications will change across a whole occupation or industry. It happens piecemeal and imperceptibly, a job at a time, a workplace at a time, driven by a spread of practices, expectations and tacit knowledge that does not follow predictable or linear patterns and is often perceptible only across a large population. Really good data on substantive complexity, if they existed, could capture the cumulative effect of such processes, but my bet is that it would show up earlier in individuals' perceptions of the changing skill-intensity of their jobs.

The other useful characteristic of skill-intensity as a common indicator of trends over time and across industries is that it is conceptually distinct from the actual technical content of the skill. It is one of the few metrics capable of supporting valid comparisons across the full range of technical content, across manual, cognitive and behavioural/interactional skills, and across levels in the skill hierarchy. This last characteristic reduces the confounding influence of social construction, since it requires workers to

compare their capability with what their own job requires, rather than with someone else's abilities. For the same reason, it minimises the risk in relying on self-report data, since informants might conceivably have an incentive to over-report the extent of their own skilfulness or that of their job, but are much less likely to see advantage in misreporting the degree of match between the two.

Note: I am not trying to suggest that a composite metric built only on these two dimensions will tell us all we need to know about changing skill requirements across the labour market. Specifically, skill-intensity is a good basis for comparing the amount and direction of change occurring over a given time period in different industries, but of little use in comparing their skilfulness at a point in time, precisely because of the insensitivity to technical content just mentioned. The only really acceptable metric would be one that combined all three dimensions. Even if this were practicable, it would be hard to define a solid rationale for combining the three into a single index; indeed, trying to do this would have wiped out one of the few really striking findings to have come out of the research reported in this paper. Nevertheless, the combination of these two dimensions provides a practical and defensible basis on which to identify trends.

The data

HILDA is Australia's first and only large-scale panel survey designed to provide a broad range of data on social and economic issues affecting the welfare of households. It shares many features with other established national surveys such as the British Household Panel Survey, the US General Social Survey and the German SOEP, suggesting that these should be investigated for their potential to support international comparative research on skilling issues. Originally designed, and still primarily intended, to serve the research needs of the Commonwealth department responsible for social security, HILDA has been progressively expanded in scope to meet the needs of other social and economic research.

At the time of writing seven annual waves of data had been released, running from 2001 to 2007. So far funding has been allocated for another five waves, taking the guaranteed run through to 2012. The sample is a panel, i.e. the same respondents are intended to remain in the sample permanently once selected. It covers around 12-13,000 households and 20,000 individuals, of whom between 12,000 and 14,000 have provided individual-level data in each wave. Of these, around 7-8,000 are employed at the time of each survey and hence eligible to answer the skill-related questions.

These questions represent only a very subsidiary element in a dataset covering more than 3,000 variables per wave. They form part of a longer sequence in which respondents are asked to rate their agreement with statements about aspects of their main job on a 7-point Likert response scale. The core set, asked over all seven waves, consists of six variables which break into two logical subsets, one referring to the skill demands of the job, the other to the degree of control or discretion which the respondent exercises over how s/he goes about the work. The variable names listed below are convenience titles adopted for the purpose of these analyses:

COMPLEX - My job is complex and difficult
NUSKILLS - My job often requires me to learn new skills
USESKILL - I use many of my skills and abilities in my current job
OWNTASK - I have a lot of freedom to decide how I do my job
HAVESAY - I have a lot of say about what happens in my job
WORKFLOW - I have a lot of freedom to decide when I do my work

Principal component analysis of the response confirms the intuitive impression that these variables fall into two reasonably reliable scales[4] covering the two dimensions of the metric. These composite scales are simple additive ones with 21 points.

From Wave 5 onwards, a supplementary set of nine variables was added, providing further detail on these issues as well as data on work-intensification. This extra detail has made it possible to construct extended scales for both primary dimensions and new sub-scales covering autonomy, control over job content and control over the organisation of working time. These new scales all show much greater reliability (Cronbach's alpha around .9 even on the 3-item scales), but are not covered in this paper because so far each covers only three waves, and most of the movement so far falls below acceptable levels of statistical significance.

Findings

Trends over the first four waves show the mean scores in decline for both indices and all but one of the constituent variables making up each scale. These trends are especially apparent for the composite scales (Figure 1), and at first sight appear to be some of the strongest evidence yet produced anywhere for an overall deskilling trend.

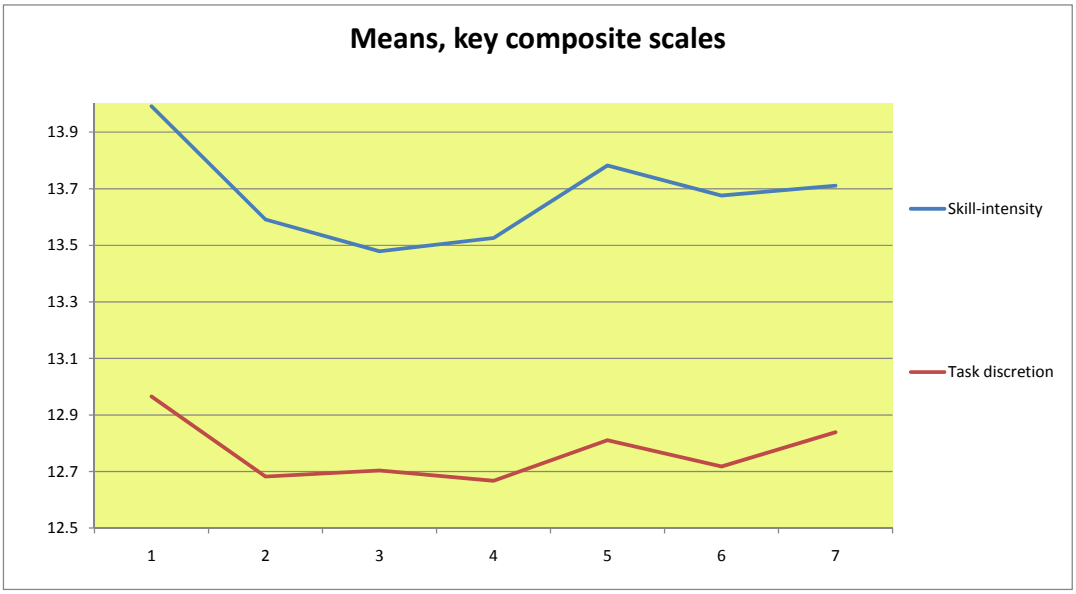

Figure 1: mean scores, composite scales, waves 1-7

On closer examination, however, the story appears less straightforward. In the first place, the movement is actually quite small, and scarcely any of the year-on-year variation reaches the 95% level of statistical significance. (The graphs visually exaggerate the scope of movement for the sake of making its direction clearer.) Indeed, were it not for the spike in scores in Wave 5, there would be no significant change in the task discretion scale between Waves 2 and 6. More specifically, the trends are not consistent but marked by a relatively steep drop between the first and second waves and an

[4] Cronbach's alpha lies a little over .7 for skill intensity and around .8 for task discretion, depending on the wave. These results are considered marginal (especially in the former case) but acceptable for 3-item scales.

almost equally sharp rise between Waves 4 and 5. This latter dropped back a little in Wave 6, but the Wave 7 results suggest it may signal the start of a new longer-term rising trend.

No explanation is immediately apparent for either of these spikes. It is tempting to suppose that the higher scores in Wave 1 are a response artefact reflecting the fact that respondents were seeing the questions for the first time and were still experimenting with their ratings. Figure 2 illustrates this effect by taking the group of respondents who gave each rating on the scale for USESKILL in the first wave and following its average score on the same question through Waves 2-7. The results (which are broadly similar for the other core questions) show that the spread of scores diminished sharply in Wave 2 and then converged much more slowly on the mean. However, the number of respondents who recorded extreme scores in Wave 1 is probably not sufficient to account by itself for the whole decline.

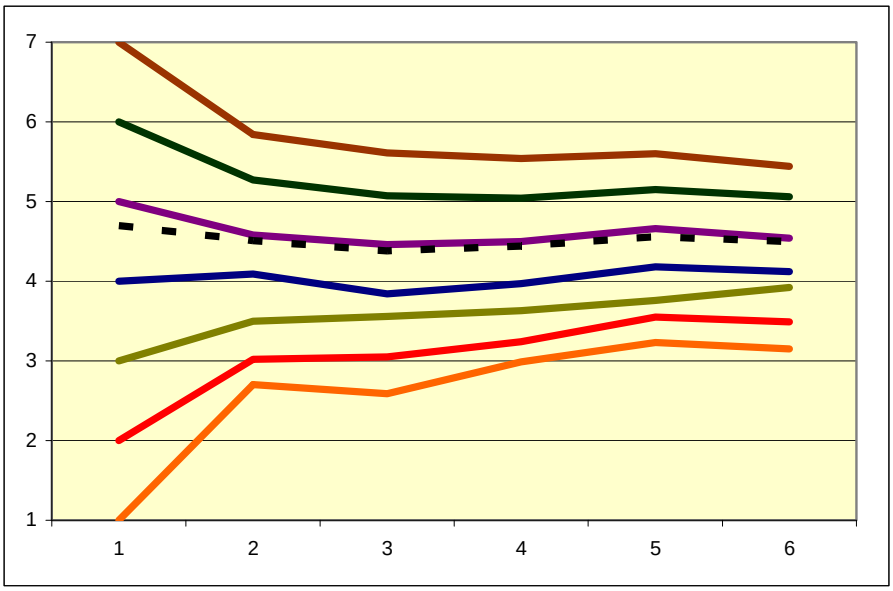

Figure 2
USESKILL: Wave 1 scores and mean scores in later waves
(dotted line: mean, all respondents in wave)

The Wave 5 spike is even harder to explain as a response effect, though it was evident across most of the of job quality variables, including some negative ones. It may be a reaction to the addition of the nine new questions, but there is no way of telling for certain.

Equally, there is no compelling reason to conclude that the high scores in Wave 1 were not genuine. The main reason it is so hard to make sense of these data is that we have no information whatever, from any source, on the trends or situation before 2001. This uncertainty is deeply frustrating, as without the sharp fall between the first two waves, the full seven-year trend on both scales would be a rising one.

To add to the uncertainty, the component variables in each scale behave differently over this period. Of the three skill-intensity variables, NUSKILLS shows the greatest variation from year to year and COMPLEX the least. COMPLEX, the lowest scoring, is also the only one to record a slight rise in the trendline over all seven waves. The scores for USESKILL appear surprisingly and consistently high,

especially when viewed against more or less comparable British evidence such as the WERS which appears to indicate much higher levels of overskilling in the UK (Mavromaras et al 2007), but this may simply be an artefact of the way the question is expressed (i.e. it is too "easy" in the terms of item response theory).

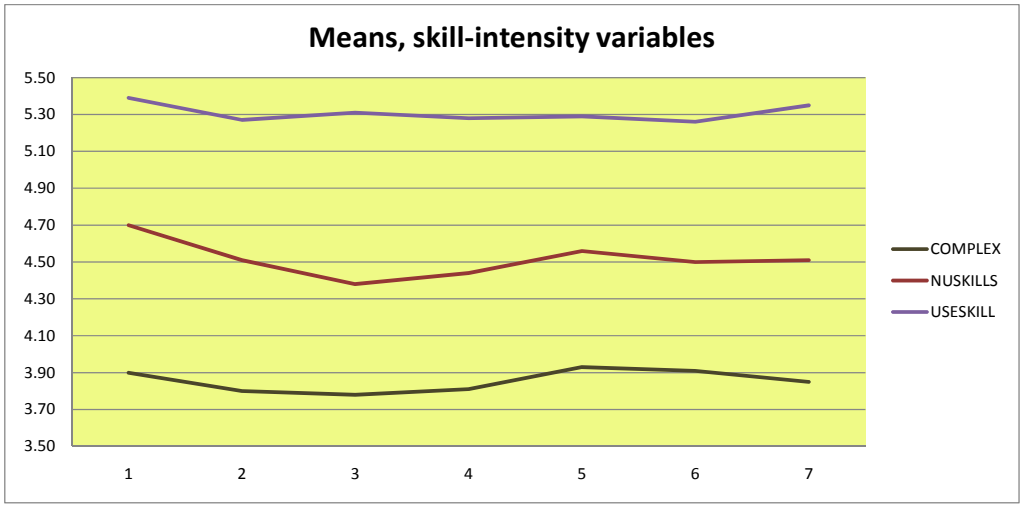

Figure 3: Components of skill-intensity, 2001-07

For task discretion it is once again the highest-scoring variable, OWNTASK, which shows least variation. Here too the lowest-scoring variable, WORKFLOW, is the only one among the six core variables whose mean score has risen steadily since Wave 2, although the amount of change is so small that it lies well within the margin of statistical error. Indeed, the overall movement in this composite scale over the six years would be negligible but for the relatively sharp drop in HAVESAY between Waves 1 and 2. It is tempting once again to dismiss this as a response effect, except that there is no obvious reason why it should have produced an artefact different from the other questions in the same scale.

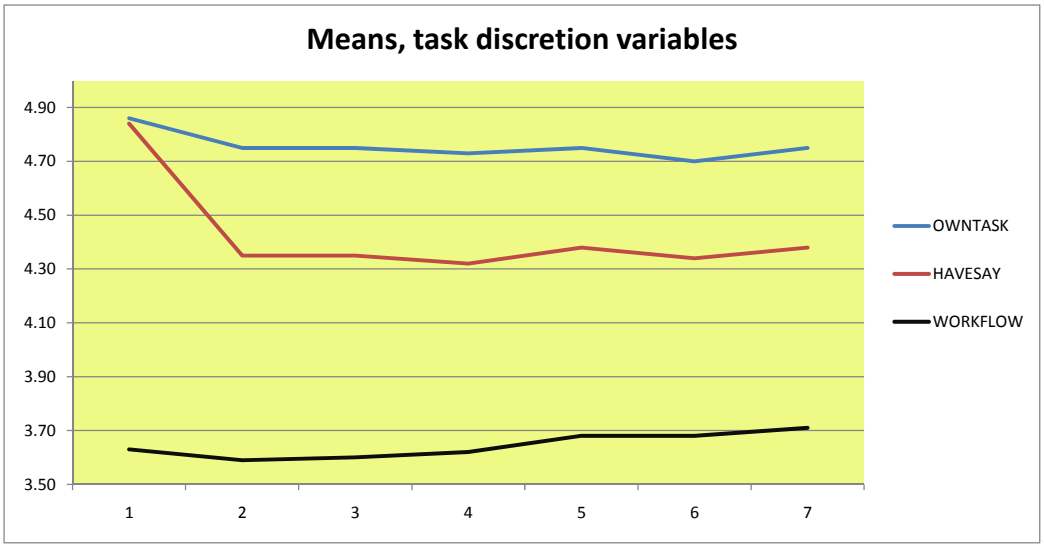

Figure 4: Components of task discretion, 2001-07

Perhaps more interesting than the overall trends, at least so far, is a discovery that came out only in the process of testing the composite scales. To explain this I need to drop for a moment into statistical technicalities.

Principal component analysis (PCA) is a way of finding out which questions in a broader set are seen by respondents to belong together as elements in common constructs. As a first step before testing the actual composite scales to be used in the research, I carried out a PCA on the full set of twelve original work quality questions. The analysis produced four such clusters, known technically as factors. In descending order of strength of association, these corresponded to task discretion, job stress, job security, and finally skill-intensity. The four factors account respectively for around 24%, 21%, 13% and 8% of the total variance. The exercise was repeated with the addition of the nine new variables in Wave 5, and this time resulted in five factors with an Eigenvalue exceeding 1 (the normal criterion used in PCA): task discretion, job stress, job security, skill-intensity and work intensification. The contribution of task discretion to variance remains virtually unchanged from earlier years, while that of skill-intensity falls to a little over 6%. Although various interpretations can be put on these findings, the most compelling is that when workers assess the quality of their job, the amount of task discretion it offers is a much more important consideration than the need or opportunity to use skill.

The robustness of the skill-intensity construct also appears much shakier in this broader analysis than the reliability testing for the actual scale would suggest. In Waves 3 and 4, indeed, there are only three factors with Eigenvalues exceeding 1, and the skill-intensity variables load on the job stress factor. COMPLEX, always the weakest item in the scale, appears particularly ambiguous, loading more strongly on job stress than on skill-intensity in Wave 5 and moderately on both in the first two. This suggests that respondents themselves see complexity as something different from the other components of skill-intensity.

From all this it can be concluded, at least tentatively, that for many Australian workers a skilful job is not seen as an unequivocal positive, since it is associated with stress and possibly work intensification. This interpretation coincides broadly with some of the findings discussed by Green (2006: Ch 4). The low ranking of the skill-intensity factor thus casts some doubt on the reliability of the scores, since if workers really pay less attention to this aspect of their job, they may well be less discriminating when asked to rate it. It is even possible that the relatively creditable mean scores achieved on two out of the three components of skill-intensity are not the good news they initially seem to be, but simply a reflection of their limited salience.

Task discretion also appears more complicated on closer inspection of the loadings. From Wave 5 onwards they are strongest for those variables which refer to individual autonomy and to control over work timing. The impression that time control may need to be considered separately from the other elements of task discretion is reinforced by the difference in trend between WORKFLOW, the one question present in all waves which relates to the time sub-dimension, and the other two items in the core scale. (The subscale for time control from Wave 5 onwards also shows a more consistent rising trend than those for autonomy and job content control, but the run of data and the range of year-on-year variation are both too small at this stage to allow confident extrapolation.) In this context the emergence of a fifth factor, again relating to control over the use of time and clearly distinguished by respondents from the other aspects of job stress, appears especially interesting. Together these findings suggest that task discretion deserves much more careful attention in future research, and that the time-related elements need further investigation in their own right.

A different aspect of task discretion has been the source of the other big surprise to come out of this research. This concerns the relation between task discretion and skill-intensity. For the full sample,

the correlation between the two scales runs a little over .2 fairly consistently over the seven waves – modest, but stronger than for most of the other expected contributory variables. However, once the sample is disaggregated by occupation or industry, some surprising discrepancies appear. Table 2 shows some of the most striking results for Wave 7.

Industry (ANZSIC 2006)	Rank, skill-intensity (out of 64)	Rank, task discretion (out of 64)
Preschool and school education	3	55
Public Order, Safety and Regulatory Services	4	54
Defence	2	45
Hospitals	5	59
Rail Transport	13	58
Finance	11	48
Coal Mining	12	47
Building Cleaning and Pest Control	63	11
Personal and Other Services	33	9
Furniture and Other Manufacturing	38	8

Table 2
Comparison of mean skill-intensity and task discretion rankings, selected industries, Wave 7

Looking at many of these industries individually, it is easy enough to find logical reasons why they should offer relatively little scope for autonomous decision-making, irrespective of their skill content. However, there are some cases, notably in mining and industries in the public sector employing large number of professionals, where it is hard to see these discrepancies as anything other than spectacularly dysfunctional, provided one accepts the existence of any link at all between the two dimensions. And even if individual discrepancies appear intuitively reasonable, their frequency suggests that the relation between the two is sufficiently problematic to put Spenner's "pragmatic hypothesis" into some doubt as a general principle.

What does it all mean?

When I first came across these data, some three years ago, a strong and counter-intuitive story seemed to be emerging. In a period when workforce skills had become one of the dominant policy obsessions in Australia, the evidence up to that point suggested that skill was getting less valued, less utilised and less developed in the average Australian workplace as time went by. Nothing of this kind had previously been demonstrated as compellingly as this, by hard data - not just in Australia, but anywhere else so far as I know.

Here as so often, the passage of time has brought with it an enforced sobriety. All it took was a couple more years of data for my trend to start looking suspiciously like a random walk. After the first seven waves of the survey, it is much harder to be certain that there is any story at all in them, never mind what that story is.

The lesson, I suggest, is not to throw out the story just yet, but to keep on analysing each new wave of data as it comes in, in the reasonable expectation that eventually a clearer story will show up (and it

could well be the original one). Remember, this is the first time anyone has tried to investigate these phenomena using annually refreshed microdata. This means we simply don't know yet how we should expect the labour market to behave, either in normal circumstances or at critical historic turning points, over such a short cycle or such a short overall timeframe; and because of that we don't know whether what we see in these data is business as usual, or part of a historically significant change. That much at least might become clearer once we have a couple more waves and can see the impact (if any) of the GFC.

It is also important to remember that the evidence I am using casts no real light on what most researchers have so far seen as the core element of skill change, substantive complexity. Even assuming the trends originally apparent in skill-intensity are indicative of the longer-term behaviour of the system, they could still imply a system lacking in dynamism rather than one lacking in skill as such. The decline evident in Waves 1-4 could be read as indicating that a majority of Australian jobs continued to require the same, perhaps reasonably high levels of skill, but that over time their occupants felt less challenged by them simply because of greater familiarity, or because their training (either before or in the course of their employment) had left them adequately prepared to respond to such change as took place in their work requirements.

So far as the story goes, the evidence for a downturn in the average skill-intensity of Australian jobs in the first five years of this century is robust enough to suggest that some kind of real trend lay behind it, even when one discounts some of the difference between the mean scores in the first two waves as a panel conditioning effect. Similarly, the pickup after 2005 seems solid enough to be genuine. In neither case can we find an easy explanation. However, at least four credible scenarios suggest themselves, all fully consistent with the known trends:

1. Both scales had shown a consistent trend (rising or falling) over previous decades but reached a plateau after around 2000. This is roughly the trend apparent in the UK Skills Surveys;
2. The variation in the seven years for which we have data fell within the normal range for the preceding decades, indicating no real change in the longer-term skill trajectory. Such a hypothesis is perfectly tenable when one considers that most of the socio-economic trends one intuitively associates with large-scale skill change – globalisation, managerialism, increasing participation in the later years of education, the transition to a post-manufacturing economy, the rollback of industrial democracy – had been in steady progress in Australia for two and even three decades by the time these data were collected, with no landmark development or turnaround in any of these areas observable over this period;
3. The higher mean scores in Wave 1 were an anomaly resulting from a response effect, and did not reflect the true situation. In this case the pattern visible after 2002 could be simply a continuation of a much longer-term rising trend;
4. The data support the commonsense expectation that skill usage becomes more efficient as the market for skilled labour tightens. It could be significant in this context that the official Skilled Vacancy Index peaked in June 2000 and declined, albeit only slightly and unevenly across occupational divisions, up to 2005 (DEWR 2005). Under this scenario, the labour market could have passed through a brief period of relaxation over the years 2002-04 when employers felt less pressure to make the most effective use of the skill pool at their disposal, which ended with evidence, and perhaps more importantly with spreading expectations, of imminent critical skill shortages from 2005 onwards.

While the last explanation would be the most satisfying (intellectually as well as in a policy sense), it must be stressed once again that the data provide no reason to prefer it over any of the others. That the same evidence can be consistent with four such disparate interpretations should serve as a warning that even the best and most current statistical evidence can provide only so much enlightenment when it is confined to such a short period. The best explanation may become clearer once a few more waves of data have left us clearer about the normal range of variation over a decade or more, especially if more robust trends have emerged by then. Pending that, it probably makes more sense to go back to qualitative research in the hope of finding at least an impressionistic explanation.

The more useful story so far, in policy terms, is that the disaggregated figures reveal the existence of several problem industries and occupations where the gap between skill requirements and the decision-making latitude available to those highly-skilled employees demands to be studied as a possible barrier to productivity. A different set of problem industries lie right at the bottom of the both scales and have ever since Wave 1, with their mean scores in many cases declining over the six years. Such industries, which mostly cluster around various aspects of food, hospitality and retailing, need to be viewed as priorities for skill upgrading – which for this purpose, remember, means not just more training but more opportunities and incentives to use skill - because of their role in important supply chains, but also because of their periodic function as sponges to soak up unemployment in a recovery. One irony is that many of the problem industries in both categories are quite strongly unionised, at least by current Australian standards, suggesting that union representation – what I call third-order industrial democracy – is not necessarily a reliable guarantee of the second order (participatory, devolved or consultative decision-making in the workplace) or of the first (individual autonomy).

Beyond that point, frustrating as it must seem to those who were hoping for social analysis, the lessons to be drawn from this research so far are of interest primarily to methodologists. To summarise:

1. Microdata are helpful for tracking skill trajectories, even if the run required to detect a genuine, sustained trend is little shorter than would be required if the surveys were taken at 4-5 year intervals like those in the UK. In the present case, the findings would look very difficult depending on whether the four-yearly slices had been taken in, say, 2000 and 2004 or 2003 and 2007.
2. Skill-intensity (despite the ambiguities attaching to COMPLEX in the present case) has proved its potential as a sensitive indicator for picking up change over very short cycles (even if much of this inevitably consists of random variation).
3. Both the skill-intensity and the task discretion indicators are effective at cancelling out most of the bias arising out of occupational prestige, and hence of social construction, as the rankings shown in Table 2 coincidentally demonstrate. The qualification is that they work for comparing the pace and direction of change in different industries and occupations, but not for comparing their relative skilfulness at a single point in time.
4. Both these indicators make it possible for meaningful data to be collected without a heavy respondent burden. In both cases a three-item scale is informative, reliable and sensitive to small changes.
5. Conversely, the relation between skill-intensity (or substantive complexity) and task discretion is much less straightforward than Spenner's original hypothesis might suggest. More useful information will come from measuring the match or gap between the two than from assuming them to be complementary and trying to combine them in a single index of skilfulness.
6. Part of the answer to this dilemma may involve treating the task discretion in each job as involving two elements, one inherent in the nature of the job (e.g. its place in the workplace

hierarchy, the practical requirement for close coordination between co-workers, the context and institutional culture in which it has to be carried out) and another, discretionary component which interacts more directly with skill. On this assumption, comparisons of change in task discretion might yield more valid and useful information if carried out within rather than between industries or occupations.

7. The time control element of task discretion appears to behave differently from those which relate to the content or manner of work. Not only does it deserve greater investigation in its own right (especially given its apparent salience to workers when they assess their jobs), but it deserves to be disaggregated more finely into its constituents. One element which could benefit from better data is the amount of control workers can exercise over the sequence in which they perform tasks. Such a variable has particular relevance to Polanyi's concept of skill, further developed by Nelson and Winter (1982: 73), which sees effective sequencing as a central generic element.
8. The model of skill change suggested earlier in this paper, whereby changes diffuse incrementally and in unpredictable patterns, implies that data on aggregates, either across the workforce or at industry/occupation level, will not necessarily will be the most effective at tracking such change as it occurs. To capture it while it its still in process, it may be more effective to examine the trajectories of representative individuals, or else of cohorts. Panel data come into their own for such purposes. HILDA is just at the stage of realising its potential for such analysis, now that sufficient waves of data are available to separate such trajectories and cohort effects from random variation and one-off effects. The future of the present research, and of research on deskilling/upskilling in general, may well lie in this kind of approach.

REFERENCES

Attewell P 1987: The Deskilling Controversy. *Work and Occupations* 14:3, 323-346

- 1990: What is skill? *Work and Occupations* 17:4, 422-448

Australia, Department of Employment, Workplace Relations and Small Business 2005: *Vacancy Report, September 2005*

Buchanan J 2006: From 'skill shortages' to decent work: the role of better skill ecosystems. Sydney, NSW Board of Vocational Education and Training

Edquist C 2004: Systems of Innovation: Perspectives and Challenges. In Fagerberg, Mowery and Nelson, 181-208

Fagerberg J, Mowery D C, Nelson R R 2004: *Oxford Handbook of Innovation*. Oxford, Oxford UP

Fagerberg J 2004: Innovation: A guide to the literature. In Fagerberg, Mowery and Nelson, 1-27

Felstead A, Gallie G, Green F 2002: *Work Skills in Britain 1986-2001*. Downloadable from www.skope.ox.ac.uk

- 2004: Job Complexity and Task Discretion: Tracking the Direction of Skills at Work in Britain. In Warhurst, Grugulis and Keep, *The Skills that Matter*, Basingstoke, Macmillan Palgrave, 148-169

Felstead A, Gallie G, Green F and Zhou, Y 2007: *Skills at Work, 1986 to 2006*. Downloadable from www.skope.ox.ac.uk

Finegold D 1999: Creating Self-sustaining, High-skill Ecosystems. *Oxford Review of Economic Policy* 15:1, 60-81

Finegold D, Soskice D 1988: The Failure of Training in Britain: Analysis and Prescription. *Oxford Review of Economic Policy* 4:3, 21-53

Form W 1987: On the Degradation of Skills. *Annual Review of Sociology* 13, 29-47

Gallie G, Felstead A, Green F 2006: Changing Patterns of Task Discretion in Britain. *Work, Employment & Society* 18:2, 243-266

Green F 2006: *Demanding Work: The Paradox of Job Quality in the Affluent Economy*. Princeton UP

Hall B H 2004: Innovation and Diffusion. In Fagerberg, Mowery and Nelson, 459-485

Hall R, Lansbury R D 2006: Skills in Australia: towards workforce development and sustainable skill ecosystems. *Journal of Industrial Relations* 48:5, 575-592

Lowry D, Molloy S, McGlennon H 2008: Future skill needs: Projections and employers' views. Adelaide, National Centre for Vocational Education Research

Malerba G 2004: Sectoral Systems: Why Innovation Differs Across Sectors. In Fagerberg, Mowery and Nelson, 380-406

Mavromaras K, McGuinness S, O'Leary N, Sloane P, Fok Y K 2007: The Problem of Overskilling in Australia and Britain. IZA Discussion Paper 3136. Bonn, Forschungsinstitut zur Zukunft der Arbeit

Nelson R, Winter S G 1982: *An Evolutionary Theory of Economic Change.* Boston, Harvard University Press

Polanyi M 1969: *Knowing and Being. Essays by Michael Polanyi*. Ed. M Grene, University of Chicago Press

Spenner K I 1983: Deciphering Prometheus: Temporal change in the skill level of work. *American Sociological Review* 48:6, 824-837

- 1985: The Upgrading and Downgrading of Occupations: Issues, Evidence and Implications for Education. *Review of Educational Research* 55:2, 125-154

- 1990: Skill: Meanings, Methods and Measures. *Work and Occupations* 17:4, 399-421

Wilson R, Hogarth T eds. 2003: *Tackling the Low Skills Equilibrium: A Review of the Issues and Some New Evidence*. University of Warwick, Institute for Employment Research