CSAE Working Paper WPS/2016-04

China's Expansion of Higher Education: the Labour Market Consequences of a Supply Shock

John Knight[1], Deng Quheng[2] and Li Shi[3]

March 2016

[1]Corresponding author. Emeritus Professor, Department of Economics, University of Oxford, and Visiting Professor, School of Economics and Business Administration, Beijing Normal University john.knight@economics.ox.ac.uk

[2] Research Fellow, Institute of Economics, Chinese Academy of Social Sciences, Beijing dengqh@cass.org.cn

[3] Professor, School of Economics and Business Administration, Beijing Normal University lishi@bnu.edu.cn

Centre for the Study of African Economies
Department of Economics · University of Oxford · Manor Road Building · Oxford OX1 3UQ
T: +44 (0)1865 271084 · F: +44 (0)1865 281447 · E: csae.enquiries@economics.ox.ac.uk · W: www.csae.ox.ac.uk

Abstract

In the decade 1998-2008 China expanded enrolment in higher education almost six-fold. For the examination of its short term labour market consequences, this unprecedentedly huge and sudden policy change might be regarded as a natural experiment. After providing a theoretical framework for analysis, the paper uses urban labour market surveys to analyse how the labour market adjusted to the supply shock**.** Three outcomes are examined: the effect of the expansion on wages, on unemployment, and on access to 'good jobs'. The shock is found to reduce relative wages, raise the unemployment rate, and reduce the proportion in good jobs, but only for the entry-year or entry-period cohort of graduates. The effect is fairly powerful for entrants, especially university rather than college graduates, but incumbent graduates are largely protected from the supply shock. An attempt is made to examine the labour market effects of the quantitative expansion on educational quality. The paper provides insight into the operation of China's labour market in recent years.

Key Words. China; cohort effects; graduate unemployment; higher education; labour market; returns to higher education

JEL Classification. I21, I23, J24, J31

1. Introduction

Beginning in 1999, China made a Great Leap Forward in higher education. In 1998 there were 3.4 million students enrolled in (regular institutions of) higher education. A decade later, in 2008, there were 20.2 million. This implies an almost six-fold growth in enrolment, averaging 19% per annum. The abruptness of the change is illustrated in Figure 1, showing enrolment from 1988 to 2008. The number of entrants to higher education was 1.1 million in 1998 and 6.1 million in 2008. Similarly, the number of graduates from higher education rose from 1.0 million in 2001 (the 1997 or 1998 intake) to 5.1 million in 2008 (the 2004 or 2005 intake).[4] How did the labour market respond to such a large shock – internationally exceptional and possibly unprecedented?

Section 2 provides the background to the policy reversal. Section 3 outlines some relevant theory of the effects of a positive supply shock of a particular type of labour on the labour market, and Section 4 briefly reviews the relevant empirical literature. Our data sources and their relevance to this issue are explained in Section 5. Section 6 conducts an analysis of the two China Household Income Project (CHIP) surveys that straddle the supply shock. Section 7 examines whether the CHIP surveys can distinguish between the direct effects of quantitative expansion and the indirect effects of that expansion on quality. Section 8 reports the empirical estimations obtained from the annual National Bureau of Statistics urban household surveys covering the years 2002-2008. Section 9 concludes and reflects.

2. A Natural Experiment?

[4] Source: National Bureau of Statistics, *China Statistical Yearbook 2009*, tables 2.5, 2.6 and 2.7. 'Regular institutions of higher education' include both university (four year degree) and college (three year degree) students.

During the first fifty years of Communist Party rule, China developed a 'flat' educational pyramid: the primary (net) enrolment ratio in 1998 (91%) was similar to that for the average of all middle income developing countries (92%), the secondary (net) enrolment ratio (50%) was a little below the average for such countries (the gross ratios being 62% and 67%), and the tertiary (gross) enrolment ratio (6%) was well below the average (12%).[5]

Figure 1. Enrolment in regular institutions of higher education in China, 1988-2009 million students

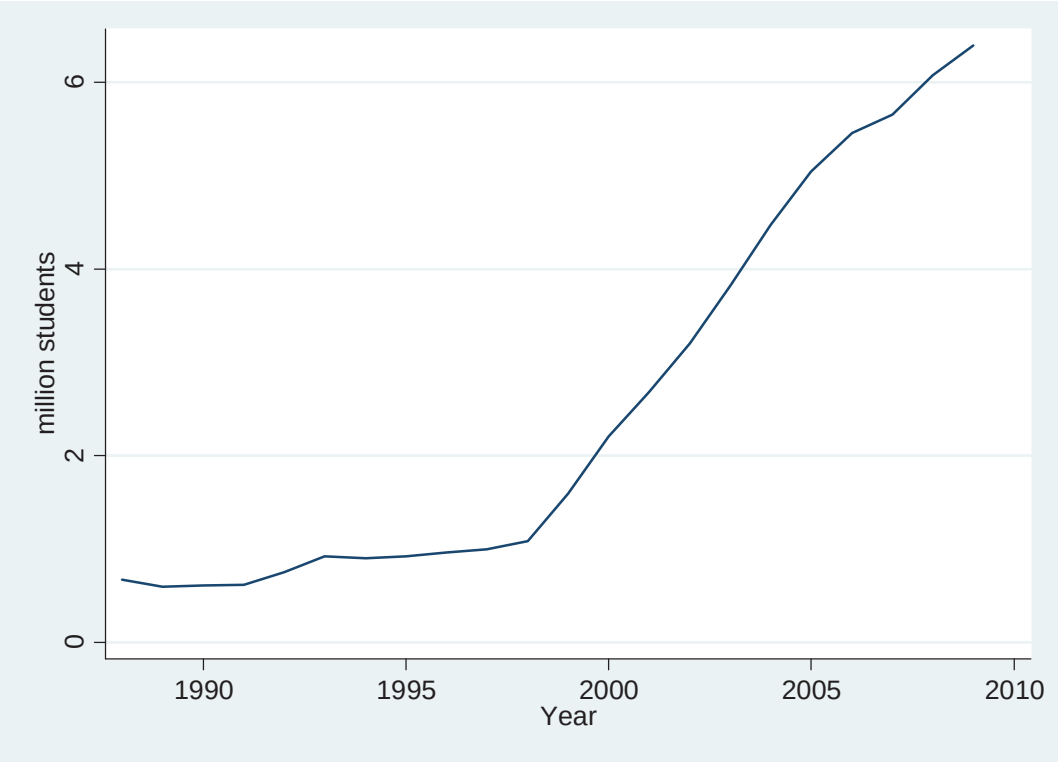

The Chinese economy grew rapidly over the decade 1988-1998 (by 9.6% per annum) and over the decade 1998-2008 (by 9.8% per annum), and the non-primary sector grew even more rapidly (by 11.2 and 10.3% per annum respectively).[6] It

[5]World Bank, *World Development Indicators 2002*, table 2.12. The net ratio (excluding over- and under-age children) is provided whenever it is available.

[6]Calculated from National Bureau of Statistics, *China Statistical Yearbook 2009*, tables 2.3 and 2.5.

would be plausible, therefore, to view the expansion of higher education as an endogenous response to the past and expected future economic growth. However, the facts suggest otherwise. There was no mention of a policy change in the Ninth Five-Year Plan (1996-2000), published in 1996. The Plan stated that student enrolment in regular institutions of higher education would grow by 3.8% per annum and reach 3.5 million in 2000.

The initial decision appears to have been taken in the light of the huge labour redundancy programme in the state-owned and urban collective enterprises that was at its height in the late-1990s. It became very difficult for young people to find employment. For instance, a national household survey conducted in 1999 found an urban unemployment rate among 17-20 year-olds in the labour market to be no less than 56%. One of the policy responses was to allow more young people to stay on in education by expanding higher education. A central government decision was suddenly taken in 1999 to raise the intake into higher education later that year by 50%.

This timing coincided with other important changes that were taking place in the labour market. In the past, government had tightly rationed but fully or mainly subsidised enrolment in higher education. It had allocated all graduates to state jobs, albeit at low wage premiums that were administratively determined. The redundancy programme in the state sector, the rapid growth of the non-state sector, and the strengthening of market forces in the labour market transformed the situation. New graduates were no longer confident of state sector jobs. The majority would have to find employment in the private sector, albeit at wage premiums that were governed by the market - but these were trending upwards. For instance, the

wage premium of higher education over high school education among urban employees was 17% in 1995 and 42% in 2002 (Knight and Song, 2008: 228)

Together with the high social value placed on education in Chinese culture, the perceived rising private rate of return to higher education ensured that the additional places were filled; the demand existed to meet the additional supply. Indeed, average tuition fees in the (exclusively state) universities could be simultaneously raised, from Y800 per annum in 1995 to Y5000 per annum in 2005. The state effectively accepted that higher education should be marketised (but not privatised).

Although the initial policy shock was a response to a short term problem, the policy evolved rapidly into one of sustained expansion of higher education. For instance, the Tenth Five-Year Plan (2001-2005), published in 2001, projected higher education enrolments to rise from 11 million in 2000 to 16 million in 2005, and the Eleventh Five-Year Plan (2006-10), published in 2006, projected enrolments to rise from 23 million in 2005 to 30 million in 2010.[7] This can be seen as a way of preparing the economy to adopt a more skill-intensive and technology-intensive form of economic development as the abundant supply of unskilled labour dries up.

Thus, the increased supply contains an endogenous element, based on past and expected future demand for graduate skills. Nevertheless, the suddenness, unexpectedness, speed and size of the supply shock means that for our purposes of analysing the short term labour market consequences, it might be viewed as creating a 'natural experiment'.

[7] *The Outline of the Eleventh Five-Year Plan for National Education Development* (http://www.moe.edu.cn/publicfiles/business/htmlfiles/moe/moe_1778/200710/27737.html).

3. Educational Expansion and the Labour Market: Theory

In a perfectly competitive labour market the effect of an expansion of higher education is straightforward. There is a reduction in the wage of all higher education graduates (call them 'graduates' for short) relative to the wages of less well educated workers. There is a uniform fall in the wage premium on higher education. General equilibrium effects follow: through substitution the fall in the wage of graduates reduces the demand for workers with whom they compete, in particular for those with whom they are closest substitutes, that is, high school leavers. In that way, the wages of high school leavers also fall but to a lesser extent.

There is a second effect. Insofar as the expansion of higher education reduces the number of high school leavers, the market wage of the latter group rises. Both effects are likely to reduce the wage difference between higher education graduates and high school leavers. However, the first effect works with a lag equal to the length of a higher education degree course, whereas the second effect works without a lag.

However, labour markets are not perfectly competitive. There are liable to be adjustment lags in any country, and protection of some workers from adverse effects. This is particularly the case in China's transitional and imperfect labour market (for instance, Knight and Song, 2005). Consider the various complications that can arise.

The immediate effect of the increase in the number of graduate entrants to the labour market is likely to be an increase in their unemployment. New graduates will search for the sorts of jobs and the starting wages that preceding cohorts secured. Gradually expectations and aspirations of the unsuccessful graduates will be adjusted downwards, and they will begin to enter 'graduate' jobs at the lower levels

of pay which clear the market for graduate entrants, and also to accept 'high school' jobs – which previously were entered only by high school leavers – at levels of pay which are either customary for those jobs or which clear the market for graduate and high school leaver entrants. How gradually this adjustment occurs will depend on several factors. These include the availability of funds for sustained job search, the degree of information which job seekers and employers possess, and graduates' expectations about how the labour market will change in the near and more distant future.

In a perfectly competitive labour market all graduates are subject to the new competition from graduate entrants. However, in practice, earlier cohorts of graduates and of high school leavers are unlikely to be significantly affected by these changes, at least in the short run. They are protected either by customary rules (against wage cuts for workers in post and against firing and replacement by undercutting workers) and also by the fact that they are likely to have received training from the employer and to have acquired firm-specific human capital, valuable to the employer.

There may be a 'rate of pay for the job', either for incumbents alone or for entrants as well. In the former case graduate entrants compete among themselves on wages and they can undercut incumbents who lack job security. In the latter case there can be competition for jobs among new entrants and unprotected incumbents, based on productivity. As graduate entrants 'filter down' into lesser jobs (defined by skill level and wage level), so they may receive the pay that is customary in those jobs. Given customary wages, insofar as graduates are preferred to high school leavers in those jobs, the high school leavers filter down into still lesser jobs, and so on down the educational hierarchy.

Knock-on effects of this sort may depend on the sensitivity of the wage to education in each job. The 'occupation-specific production function' (OPF) (Knight, 1979), relating productivity in the job to education, is likely to rise with education but to encounter diminishing returns to education. The OPF is likely to be more concave (bowed) than is the relationship between productivity and education across all occupations. For instance, a certain level of education may be required for a job to be done at all, and beyond a certain, greater, level of education no improvement in productivity occurs. The OPF may imply that employers are indifferent between employing high school leavers at one wage and graduates at a higher wage reflecting their higher average productivity in the occupation.

These considerations provide a loose theoretical framework and suggest some testable hypotheses for analysing the effects on the labour market of China's Great Leap Forward in higher education.

4. Educational Expansion and the Labour Market: Evidence

The consequences of higher educational expansion were examined in Freeman (1976). College education had expanded rapidly in the United States during the prosperous times of the 1950s and 1960s, when enrolments tripled. However, the conditions facing labour market entrants with college degrees worsened sharply as the economy slowed down in the 1970s. He showed that, although the proportion of managerial and professional jobs had risen over the period 1952-74, the ratio of managerial and professional jobs to college graduates had fallen. The ratio of college starting salaries to average salaries declined over the same period and the unemployment rate of recent college graduates rose substantially.

Knight and Sabot (1990) examined the effects of secondary school expansion in the labour market in East Africa. They found that the wage structure in the market sector was compressed by the expansion, and that this occurred in two ways. Secondary school leavers 'filtered down' into lower-skill, lower-wage occupations, and their pay within occupations was reduced. There were powerful cohort effects: both the occupational distribution and the intra-occupation wages of recent secondary school leavers were inferior to those of earlier cohorts.

China has had a previous shock to its labour market: the result of the redundancy policy - at its height in the late 1990s - that led to millions of state-owned enterprise (SOE) employees being retrenched. Appleton et al. (2002) found from a survey conducted in 2000 that 11% of urban workers had been retrenched and that 53% of these remained unemployed. Appleton et al. (2004) used the same data to compare non-retrenched and re-employed urban workers. The authors inferred from the wage premium of non-retrenched workers that the labour market was segmented according to the openness to market competition. Whereas many non-retrenched workers, being in internal labour markets, enjoyed wages above market levels, many retrenched workers faced strong competition for jobs.

Knight and Li (2006) examined the employment of retrenched workers in more detail using the same survey. The average completed spell of unemployment was 19 months and the average projected 'expected' spell 47 months. There are reasons why unemployment duration and the re-employment wage can be positively related: those holding out for a high wage are likely to wait longer. However, a strong negative relationship was found. Various explanations are possible: loss of unemployment support, loss of human capital, the labour market becoming tougher,

or expected to become tougher, over time, and employer screening against the long term unemployed. China's displaced workers had been thrown onto a tough labour market.

There are a few studies of the recent labour market for college graduates in China, but none has examined the hypotheses of this paper in any detail. Li and Xing (2010) used the 1% Population Survey of 2005 to investigate how the expansion had affected access for different sorts of people. In addition, however, they used variation in age and province to show that the expansion reduced the college wage premium of the expansion cohorts and increased their unemployment.

Chi et al. (2012) analysed the Chinese Urban Household Survey over the years 1989-2009. Examining starting earnings and jobs of different entrance-year cohorts and the subsequent growth of earnings over the next few years, they found that the wages and employment of young workers varied with market conditions at the time of labour market entry but that these cohort effects diminished within three years. The authors did not concentrate on higher education and they did not distinguish between the first half of their period (before the great expansion of higher education) and the second. Their results are nevertheless at odds with the findings of (pre-expansion) labour market segmentation between incumbents and job-seekers referred to above. However, the urban labour market has marketised rapidly in recent years, and it has become more competitive in the new century.

5. The Data

We use two data sources. Both are needed because the weaknesses of one are the strengths of the other. First, we draw on two cross-section national household surveys of the China Household Income Project (CHIP): one (CHIP 2002) conducted early in 2003 and relating mainly to 2002, and the other (CHIP 2007) conducted early in 2008 and relating mainly to 2007. The advantages of these surveys for our purposes are that they are nationally representative, they contain rich research material – detailed socioeconomic information on households, their members and their communities - and that they span the period of the labour supply shock. The surveys contain good information on education [including length of degree course, relative performance at school, and (in 2007) mark in the national college entrance examination], on unemployment, and on wages.

The standard university first degree in China entails a four-year course and the standard college degree entails a three-year course, both starting in the autumn of each year. Thus, even the freshest graduates captured in the 2002 survey had enrolled prior to 1999, the year in which enrolment began to expand rapidly. By contrast, the 2007 survey contained those who had graduated in the five years 2003-2007. It is therefore possible to compare this cohort with the cohort who graduated in the five years 1998-2002. The two surveys are five years apart. In order to estimate the effect of the expansion on the preceding five-year cohort, it is possible to use the 2007 survey to compare the cohort that had graduated in the years 1998-2002 with the corresponding cohort in the 2002 survey, i.e. those who had graduated in the five-year period 1993-1997. Thus, the surveys lend themselves to difference-in-difference cohort analysis.

Full accounts of the 2002 and 2007 surveys are to be found in the two edited volumes that reported their main results (Gustafsson et al., 2009 and Li et al., 2013,

respectively). Briefly, and most relevantly, the urban survey of 2002 contains 12 provinces, 70 cities, 6,800 households and 20,700 individuals. The corresponding figures for 2007 are 16 provinces (including the original 12), 300 cities, 10,000 households and 29,300 individuals. As one third of all households are dropped by the NBS each year, the two surveys do not contain any of the same households. Both surveys are representative subsamples of the NBS national household samples. For instance, CHIP provinces and cities are chosen to be nationally representative. Households in the NBS sample maintain written records. The CHIP questionnaires provided additional information with research hypotheses in mind.

In principle it would be possible to combine the three subsamples – registered urban residents, registered rural residents, and rural-urban migrants (recorded in urban areas but with rural registration) - and to analyse all three components of the labour market together. However, there are several reasons why we choose not to do so. First, the great majority of higher education graduates are employed in urban areas. Secondly, all higher education graduates born in rural China are entitled to obtain an urban *hukou* (residence registration), whereas other rural people are generally not eligible. The social and economic privileges of possessing an urban *hukou* and the better opportunities for graduate employment in cities ensure that almost all rural graduates choose to exercise their right. Thus graduates among the rural-urban migrants are likely to waste no time in transferring to urban *hukou* status; accordingly, they appear in the urban household survey and not in the rural-urban migrant survey. Higher education graduates account for 16.9% of employees in the 2002 urban sample, 0.6% in the rural-urban migrant sample, and 0.4% in the rural sample; the corresponding figures for the most recent five-year entry cohort in the 2007 survey are 22.7%, 1.7% and 2.1% respectively. Thirdly, the institutional

differences in the labour markets for the three groups imply that separate analyses would be more reliable (Knight and Song, 2005, chs. 5-8). As the urban labour market for higher education graduates is by far the most important of the three, the analysis is confined to the urban *hukou* sample.

Our other data source is the urban samples of the NBS annual representative household survey for the years from 2002 to 2008. These samples have the advantage of being larger than the CHIP subsamples (containing 43,500 individuals in 2002 and 39,000 individuals in 2008)and of providing seven years of data, but the disadvantage of containing less information. The NBS survey is generally not a panel as about a third of households are rotated each year. The urban survey contains information on individuals: educational level, age, gender, ethnicity, *hukou* status, labour force participation, employment status, year in which started employment, earned income, hours worked, occupation, sector, and ownership of employer. It is therefore possible to analyse the determinants of the wage, of unemployment, and of being in a 'good' job.

6. Analysis of the CHIP Surveys

In the 2002 (urban) CHIP survey 25% of individuals aged 16 or above had college education or above. In the ensuing five years this had risen to 29%. For the age group 22-26 the corresponding figures were 49% and 61% respectively. In the case of university education (excluding the three-year colleges) the equivalent proportions for all adults were 8% and 13%, and for the age group 22-26, 18% and 23%.The rapid expansion of higher education over these five years is vividly apparent in the two CHIP samples.

We initially conduct a standard earnings function analysis using the two urban samples, with education levels and also a full set of conditioning variables on the right-hand side. Our objective is to discover whether the labour market performs in line with standard competitive labour market theory. The coefficient of interest is that on higher education, with completed high school as the omitted education category; this measures the wage premium on higher education. Over the five-year period the share of higher education graduates in total employment rose from 33% to 41%, that of university graduates rose from 10% to 17%, and that of high school graduates fell from 41% to 36%. Thus, the ratio of employment of higher education graduates to that of high school graduates rose from 0.81 to 1.14, i.e. by over 40%. The expectation is that the higher education premium would be depressed by market forces.

An earnings function (not reported) was estimated for the sample of higher education graduates and high school leavers and with the conventional explanatory variables (experience and its square, gender, ethnicity, and region) in 2002 and 2007. All the coefficients other than ethnicity were statistically significant and had the expected sign. The coefficient on higher education fell only a little, changing from 0.37to 0.34 (both significant at the 1% level) over the five years of rapid expansion in the number of college and university graduates in the labour market.

There are at least five possible reasons why the wage premium was relatively insensitive to the supply shock. The first is that the urban economy grew fast enough to absorb most of the entrants, although the much faster percentage growth of graduate numbers belies this explanation. Two other possibilities are bias in technological progress or a changing pattern of demand, either of which could increase the demand for higher education graduates relative to high school leavers. The testing of these hypotheses would require separate research papers. There is a fourth possible

explanation, which can be tested. It is that the supply shock is selective in its effects: although the effects do not show up in the labour market as a whole, they nevertheless exist but are confined to the most recent cohort or cohorts. The fifth hypothesis is that the wage premium is more sensitive to the expansion of university enrolment, because universities, more than universities and colleges together, bore the brunt of the educational expansion.

Similarly, a comparison of the overall unemployment rate among higher education graduates in the urban sample in 2002 and 2007 does not suggest a large supply shock: the rate was 4.6% in 2002 and 6.0% in 2007. It is true that the unemployment rate among university graduates increased by more percentage points: from 2.7 to 5.0%. Again, however, the reason might be found in the selective nature of unemployment, with the incidence being highest for the most recent cohort or cohorts.

Our first hypothesis is that the effect of the supply shock was to concentrate the fall in wages on the entry cohort or cohorts. Table 1 provides weak support in the case of higher education graduates. The table shows the ratio of higher education graduates' to high school leavers' average hourly wage by five-year age group. For the entry cohort aged 22-26, the ratio fell (by -0.09). There was a rise (+0.23) for the age group 27-31 and a smaller rise (+0.14) for the age group 32-36. Those aged 22-26 in 2002 were aged 27-31 in 2007. Their ratio rose over time (+0.20) and the ratio for the age group 27-31 in 2002 also rose as this cohort aged to 32-36 in 2007 (+0.14). It is only the entry cohort that had a lower

Table 1. The ratio of average hourly wage of higher education graduates to high school leavers by five-year age cohort, 2002 and 2007

Age cohort	2002	2007	2007 minus 2002	2007 minus 2002 five years younger

22-26	1.46	1.37	-0.09	
27-31	1.43	1.66	0.23	0.20
32-36	1.43	1.57	0.14	0.14

wage ratio in 2007 than in 2002. This suggests a cohort effect. The same exercise conducted for university graduates produces stronger support for our hypothesis (Table 2). The wage ratio for the age-group 22-26 fell by 0.37, whereas the ratio for the age group 27-31 rose by 0.11 and that for the age group 32-36 by 0.20. The table also shows that the ratio for those who had been 22-26 in 2002 had fallen by 0.02 by 2007, when they were aged 27-31, whereas the ratio over the same period for the age cohort immediately above them rose by 0.09. Cohort effects were evident, but the only falls in the ratio were among new entrants and among the 2002 new entrants as they aged five years.

Table 2. The ratio of average hourly wage of university graduates to high school leavers by five-year age cohort, 2002 and 2007

Age cohort	2002	2007	2007 minus 2002	2007 minus 2002 five years younger
22-26	1.89	1.52	-0.37	
27-31	1.76	1.87	0.11	-0.02
32-36	1.65	1.85	0.20	0.09

Our second hypothesis is that there were cohort effects also in the unemployment rate. Thus, the unemployment rate of recent graduates should be higher in 2007 than in

2002, and the rise should be greater than for older graduates. We also hypothesise that the supply shock between 2002 and 2007 reduced the unemployment rate of recent high school leavers because fewer of those completing high school left education at that stage, and the leavers did not have to compete with their peers who were now in higher education.

Table 3. Percentage unemployment rate of higher education graduates, university graduates and high school leavers by five-year age cohort, 2002 and 2007

Age cohort:	22-26	27-31
	2002	
Higher education graduates	9.5	6.6
University graduates	3.0	3.8
Leavers	26.4	18.7
	2007	
Higher education graduates	12.4	4.5
University graduates	9.0	2.4
Leavers	27.6	12.5
	2007 minus 2002	
Higher education graduates	2.9	-2.1
University graduates	6.0	-1.4
Leavers	1.2	-6.2
	2007 minus 2002 five years younger	
Higher education graduates		-5.0

University graduates	-0.6
Leavers	-13.9

Table 3 shows that the unemployment rate of the entry cohort (aged 22-26) of university graduates rose, from 3.0% in 2002 to 9.0% in 2007. For the entry cohort of higher education graduates there was a small rise, from 9.5% to 12.4% over the five years. In both 2002 (26.4%) and 2007 (27.6%) the unemployment rate of young high school leavers was much above that of the graduates. This high rate might, in 2002, reflect the effects of the powerful redundancy policies in the late 1990s and early 2000s and, in 2007, the preference given by recruiters to the, now more numerous, higher education graduates.

The difference-in-difference analysis shows the tendency for the unemployment rate to fall as length of time in the labour market rises, and more so for the cohort aged 22-26 in 2002 than for the cohort that was five years older. In summary, Table 3 provides some evidence in support of our hypothesis: for the cohort aged 22-26, the university graduate and higher education graduate unemployment rates rose, by 6.0 and 2.9 percentage points respectively, whereas the cohort aged 27-31 experienced a fall in unemployment. However, the unemployment rate of the entry cohort of high school leavers did not show the hypothesised fall.

Our third hypothesis is that, as a result of the supply shock, university graduates, higher education graduates and young high school leavers 'filtered down' into lesser jobs ranked by average occupational wage. The most useful distinction is between managerial and professional occupations (which together can be regarded as 'good jobs') and the rest: we classify occupations accordingly.

Table 4 provides the relevant evidence. Between 2002 and 2007 the proportion of higher education graduates employed in good jobs fell in both age cohorts. The fall (-10.1 percentage points) was large for the new entrants than for the age group five years older (-3.3 percentage points). It was more marked for university graduates, being -18.9 and -15.4 percentage points for the age groups 22-26 and 27-31respectively, The heading

Table 4. Percentage occupational distribution of higher education graduates, university graduates, and high school leavers by five-year age cohort, 2002 and 2007

Age cohort:	22-26	27-31
Higher education graduates		
2002		
Good jobs	42.1	47.7
Other	57.9	52.3
2007		
Good jobs	32.0	44.4
Other	68.0	55.6
2007 minus 2002		
Good jobs	-10.1	-3.3
Other	10.1	3.3
2007 minus 2002 five years younger		
Good jobs		2.3
Other		-2.3

University graduates

2002		
Good jobs	59.6	64.6
Other	40.4	35.4
2007		
Good jobs	41.1	49.2
Other	58.9	50.8
2007 minus 2002		
Good jobs	-18.5	-15.4
Other	18.5	15.4
2007 minus 2002 five years younger		
Good jobs		-10.4
Other		10.4

High school leavers

2002		
Good jobs	8.3	6.2
Other	91.7	93.8
2007		
Good jobs	7.6	17.3
Other	92.4	82.7
2007 minus 2002		
Good jobs	-0.7	11.1
Other	0.7	-11.1

2007 minus 2002 five years younger	
Good jobs	9.0
Other	-9.0

'2007 minus 2002 five years younger' can be expected to reflect progress up the job ladder as workers acquire experience and skills. It records a rise for higher education graduates (by 2.3 percentage points) but a fall in the proportion of university graduates in good jobs (by 10.4 percentage points). No more than 8.3% of high school leavers aged 22-26 were in good jobs in 2002. Over the next five years this proportion fells lightly, to 7.6%, but it actually rose for older cohort.

In summary, filtering down into lesser occupations was powerful among higher education graduates, and especially among university graduates, over the period 2002-2007, and it was sharpest, but not confined to, the younger cohort. By contrast, there is no real evidence of filtering down among young high school leavers: it appears either that they were not displaced down the job ladder by graduates accepting jobs previously filled by high school leavers or that they benefited from the labour market absence of their peers who had continued into higher education.

7. The CHIP Surveys and the Quality of Education

The rapid expansion of enrolment in higher education is likely to have produced two adverse effects on quality. If so, the fall in quality, instead of the direct rise in quantity, might account for some of the observed changes in the wages, unemployment or occupational attainment of higher education graduates. First, rapid expansion of educational institutions may have caused the quality of teaching and learning to suffer. Second, the average quality (that is, 'natural ability') of students may have declined as a

higher proportion of 18-year-olds went into higher education. It is desirable to standardise for quality in the analysis in order to isolate the effect of the change in quantity on its own. The CHIP surveys permit some standardisation but, because the two surveys contained different quality questions, only limited progress can be made.

Both surveys asked about quintile performance in secondary school. Almost all the relevant respondents reported being in the first three quintiles. Table 6 shows that in

Table 5. Quintile performance in secondary school: percentage distributions and wage coefficients

	Percentage in first quintile	Coefficient of first quintile
Higher education graduates		
All ages		
2002	18	0.164***
2007	24	0.137***
Age group 22-26		
2007	20	0.198**
2007	19	0.196***
University graduates		
All ages		
2002	44	0.108**
2007	32	0.097**
Age group 22-26		
2002	39	0.319**

2007	31	0.255***
High school leavers		
All ages		
2002	10	0.164***
2007	11	0.067
Age group 22-27		
2002	6	0.228
2007	8	0.048

Notes: The omitted category in the wage function is all four quintiles below the first. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively.

2002 28% of higher education graduates, 44% of university graduates, and 10% of high school leavers reported having been in the top quintile. There was only minor change over the five years 2002-2007: the percentage of higher education and university graduates in those quintiles fell by 4% and by 12% respectively, and the percentage of high school graduates was stable. Thus, the only notable sign of fall in quality was among university graduates.

With quintiles 2-5 as the omitted category, the conditional wage coefficient on quintile 1 fell in each case over the five years (Table 5). Performance at high school (representing pupil quality since respondents were likely to report ranking within their school) became less important for wages at each level. Our hypothesis was that as higher education, and especially university education, expanded, so employers placed more value on the quality, as opposed to the educational credentials, of their employees, but the evidence suggests otherwise. However, the results might be different for the age-

cohort 22-26, which had entered employment in the previous five years (also reported in the table). The main compositional effect was that the proportion of university graduates who reported having been in the first quintile fell by 8 percentage points. Again, however, the wage coefficients did not rise over the five years. The wage premium on good school performance was higher for new entrants than for all ages. That was true of both higher education graduates and university graduates but especially of the latter. This result suggests that, in setting wages, employers paid more attention to the educational quality of their recent recruits than of their established employees.

Table 6. The wage coefficient of higher education graduates and university graduates, with and without standardisation for quintile performance at school, 2002 and 2007

	2002	2007
Sample of higher education graduates and high school leavers		
Coefficient on higher education		
All ages		
Without quintile performance	0.373***	0.344***
With quintile performance	0.344***	0.330***
Age group 22-26		
Without quintile performance	0.292***	0.346***
With quintile performance	0.268***	0.342***
Sample of university graduates and high school leavers		
Coefficient on university education		
All ages		

Without quintile performance	0.565***	0.519***
With quintile performance	0.476***	0.458***
Age group 22-26		
Without quintile performance	0.573***	0.494***
With quintile performance	0.478***	0.468***

Table 6 examines the conditional wage premium on higher education and on university education (both relative to high school education) in 2002 and 2007. The estimates are made both with and without standardisation for quintile performance at secondary school. In each case standardising for quintile performance reduces the wage premium, but not by much. The wage premium is lower in 2007 than in 2002, as expected, but with one exception: higher education graduates aged 22-26. However, in each case the fall is lower (or the rise higher) when there is conditioning on quintile performance. These results suggest that part of the wage premium is due to the superior quintile performance of those with more education, and that the fall in the wage premium over time is

somewhat offset by the effect of quintile performance. Unlike entrants with university education – who bore the brunt of the expansion – entrants with higher education gained in wage premium over the five years, but it is their low wage premium in 2002, before the great educational expansion could have an effect on the labour market, that produces the anomaly.

Unlike the 2007 CHIP survey, the 2002 survey provided information on whether the respondent had attended a 'key' secondary school, and on the 'rank' of university attended (grouped into three ranks, rank 1 being highest and rank 3 being lowest).

Attendance at a key secondary school raised the conditional wage (relative to high school leavers) of higher education graduates (0.102***) and of university graduates (0.133***). Attendance at a key school represents a combination of pupil quality and school quality. University rank affected university graduates' pay in 2002 only if the lowest rank of university had been attended (-0.122**). Unfortunately it was not possible to test the hypothesis that the rank of university attended had a greater effect on the wages of university graduates in 2007, after the great expansion of university education required employers to be more selective in their recruitment and remuneration policies.

The 2007 CHIP survey contained some information that the 2002 survey lacked, in particular the *gaokao* score (crucial to the admission to institutions of higher education and their quality). It is not sensible to calculate measures of central tendency and dispersion of the *gaokao* score because the *gaokao* scoring system varied across years and provinces. The *gaokao* score can affect wages because it is a measure of individual ability and achievement (if it is known by employers) or a proxy for these productive qualities (if it is not known by employers). The *gaokao* score can also raise wages because it provides access to good institutions of higher education and so contributes to the amount of human capital acquired. The pure *gaokao* score is relevant if the effect on wages operates through individual qualities, but the *gaokao* score conditioned on province and year (so correcting for different admissions criteria) is relevant if the effect on wages operates through access to good institutions.

Table 7 reports the conditional effect of the *gaokao* score on wages. One point more in the *gaokao* score raised the wages of higher education graduates by a statistically significant amount and of university graduates by almost the same amount. Examining the group aged 22-26, it is notable that each significant coefficient on *gaokao* is higher for the entry cohort than for employees as a whole. This is consistent with the notion that individual ability and

human capital acquired in university or college are becoming more important in wage setting given the much enlarged field of young graduates. It is unfortunate that we cannot directly test the hypothesis that the effect of the *gaokao* score on wages increased over the five years in which higher education expanded and so quality became more important to employers. We note in the table that the coefficient on the *gaokao* score is invariably higher, although not generally significantly so, for corrected than for pure *gaokao*. This suggests that *gaokao* makes some contribution to wages through its allocative function.

Table 7. Coefficients on gaokao score with and without standardisation for province and year of test; coefficients on higher education and on university education in wage functions for samples of higher education plus high school education and of university education plus high school education, with and without standardisation for gaokao score, 2007

	All ages	Age-group 22-26
Samples by educational level		
Coefficient on *gaokao* score:		
Standardising for province and year of test		
Higher education	0.0015***	0.0019***
University education	0.0014***	0.0027***
Not standardising for province and year of test		
Higher education	0.0013***	0.0014***
University education	0.0013***	0.0019***
Sample of higher education graduates and high school graduates		
Coefficients in wage functions:		
With gaokao variable		

Higher education	0.238***	0.403***
Gaokao score	0.0014***	0.019***
Without gaokao variable		
Higher education	0.344***	0.537***
Sample of university education and high school graduates		
With gaokao variable		
University education	0.334***	0.503***
Gaokao score	0.0014***	0.0019***
Without gaokao variable		
University education	0.477***	0.495***

Notes: ***, **, * denote statistical significance at the 1, 5 and 10% levels respectively. All the coefficients are estimated using the standard wage function described in Table 5.

Table 7 also shows how standardisation for the *gaokao* score affected the conditional wage premium on higher education and on university education (again in relation to high school education) in 2007. The premium is shown with and without standardisation for *gaokao*. The table shows that the coefficient on the *gaokao* score is significantly positive in each case and that its introduction reduces the coefficient on higher education by 30 % in the case of higher education graduates and by no less than 40% in the case of university graduates. In 2007 the *gaokao* score, and its consequences, were important for wage success. The table also reports equivalent estimates for the age-group 22-26. In each case the coefficients on the wage premium for higher or university education and for the *gaokao* score are higher for new entrants than for established employees. This is consistent with the notion of employers paying more attention to the educational quality of their new recruits as the quality of graduates widened with the expansion of enrolments.

Table 8 reports the wage premia on higher education and on university education by performance quintile. There is an interesting contrast between the two years. For all ages, in 2002 there is no difference in the wage premia according to quintile. However, in 2007 the coefficient is far greater for those in quintile 1 or 2 than for those in quintile 3. Whereas the quintile performance did not appear to matter to the wage premia in 2002,

Table 8. Wage premium on higher education for sample of higher education graduates and high school leavers. and on university education for sample of university graduates and high school leavers, by performance quintile, 2002 and 2007

	2002	2007
Sample of higher education graduates and high school leavers		
Higher education graduate coefficient		
All ages		
For performance quintile 1 or 2	0.336***	0.383***
For performance quintile 3	0.339***	0.203***
Age-group 22-26		
For performance quintile 1 or 2	0.217***	0.305***
For performance quintile 3	0.334***	0.404***
Sample of university graduates and high school leavers		
University graduate coefficient		
All ages		
For performance quintile 1 or 2	0.513***	0.496***
For performance quintile 3	0.514***	0.295***
Age-group 22-26		

For performance quintile 1 or 2	0.470***	0.487***
For performance quintile 3	0.401**	0.504***

Notes: The numbers in performance quintiles 4 and 5 are very small.

it did so in 2007, when the numbers with higher education had expanded greatly. This is consistent with employers then valuing college and university graduates more highly if they possessed higher natural ability or, through that, had attended better institutions. If this neat story is correct, we might expect the effect to be accentuated among new entrants. However, that hypothesis is belied by the evidence for the age-group 22-26, also shown in the table. There is no equivalent divergence in the coefficients on education between 2002 and 2007. Indeed, at both educational levels the coefficient is higher for quintile 3 than for quintile 1 or 2. This might be because, within quintile 3, remaining high school leavers were in other respects of lower quality and were now less well paid.

8. Analysis of the NBS Surveys

The NBS surveys – described in broad terms in Section 5 - permit a more detailed and complex analysis and cover the years 2002-2008. Unfortunately there is no direct information of the year in which an individual entered the labour market. We know the year of first employment but this might follow a period of unemployment. Accordingly we make the (normally correct) assumption that higher graduates entered the labour market at age 22. With these data on educational level, year of labour market entry and current age, for each survey year we can link these variables to our three variables of interest: the real or relative wage of the individual, whether they are unemployed, and whether they are employed in a 'good job'. In this way it is possible to track the effect of educational expansion on the labour market from one year to another. However, the use

of ‘educational level-entry year-age’ cells in some cases creates a problem of small numbers, particularly serious for the university graduates. For this reason we concentrate on the (more numerous) higher education graduates.

There is a serious discontinuity in the series for unemployment rate and in the series for the proportion in good jobs. This can be observed in the unemployment rate for all workers in the sample, which falls from 11.8% in 2002 to 5.8% in 2008. However, 91% of this fall takes place between 2006 and 2007. Similarly, the proportion of all employees in good jobs rises from 24.0% in 2002 to 35.1%in 2008, but 75% of this rise occurs between 2006 and 2007. In principle, the discontinuities could be the result of a change in sampling or a change in definition. However, the fact that there is no discontinuity in the wage series suggests that the problem is one of definition. Accordingly, we analyse the unemployment rate and the proportion in good jobs only for the period 2002-2006, whereas the wage analysis includes 2007 and 2008 as well.

8.1 The effect on wages

Consider first the effect of educational expansion on wages. We use the variable ‘income from employment’ and deflate it by the urban consumer price index to measure the real wage in each year. The results are best understood by means of Figure 2, which shows real earnings of higher education graduates for each year of entry cohort and age. The entry wage remains fairly constant over the period, being lowest for the 2003, 2005 and 2007 entry cohorts and highest for the 2006 and 2008 entry cohorts. The 2003 entry cohorts do badly at age 23 and 24 and the 2004 entry cohort at age 23. Beyond entry cohort 2003, the wage of those aged 24 and over tends to rise with the year of their entry. For each entry cohort, there is normally an increase the wage as that cohort grows older. Thus, it is the rough stagnation of the entry wage that appears to be the

main effect of the higher education expansion on wages – a stagnation that does not continue as the entrants grow older.

However, this conclusion is misleading. The average real earnings of the NBS samples more than doubled between 2002 and 2008, with their growth averaging 10.5 per cent

Figure 2. Real earnings of higher education graduates by age for each entry cohort

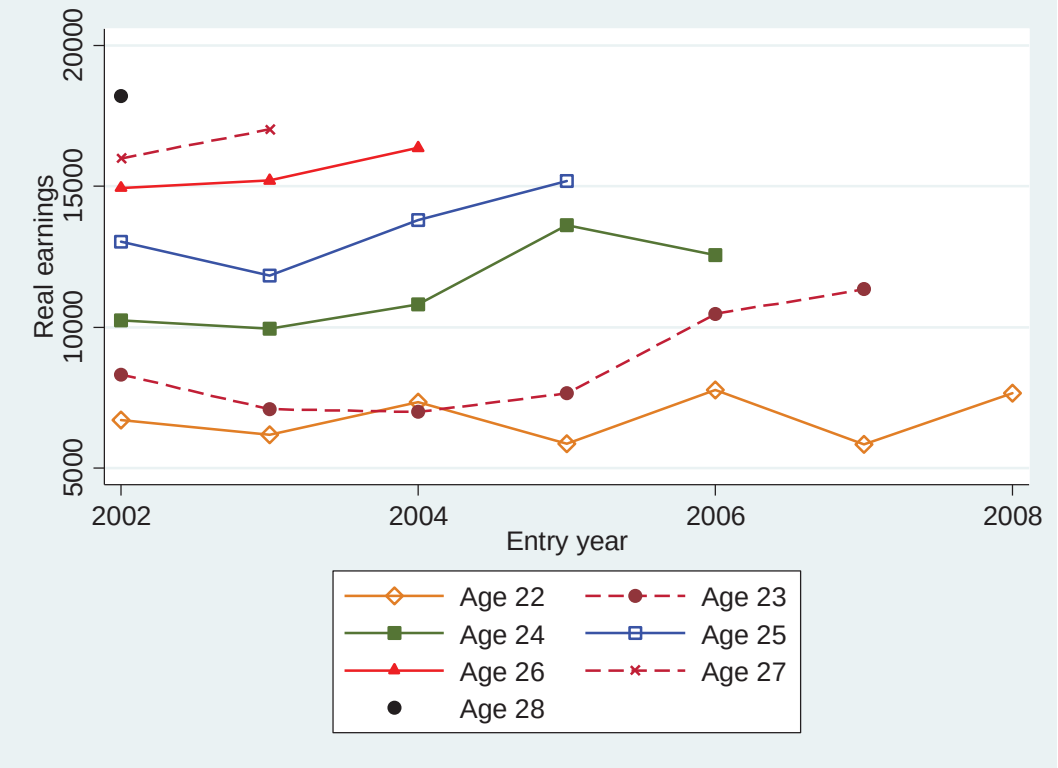

per annum. There are several possible reasons for this rapid increase: emerging shortage of labour, although this probably became a general phenomenon only after

2009 (Knight et al., 2012); the growth of minimum wages; rising labour efficiency associated with physical and human capital accumulation, reallocation of labour to more productive activities, and economic reforms. Most workers should have been able to benefit from the greater labour efficiency that either caused or was caused by the rising wages, including the subjects of our study.

It is therefore relevant to examine the changing relationship between the real wages of graduates entering the labour market and real wages in the labour market as a whole. Treating the year 2002 as the pre-expansion relationship (and so not deflated), we deflate the real wages of graduates by the proportionate increase in real wages in the labour market as a whole. If the deflated value of the graduate real wage falls, this indicates that the graduate real wage has fallen relative to real wages overall; and if it is more negative in the following year, this indicates that the relative wage has fallen again.

Figure 3. Relative real earnings of higher education graduates, deflated by the real earnings of all employees, by age for each entry cohort

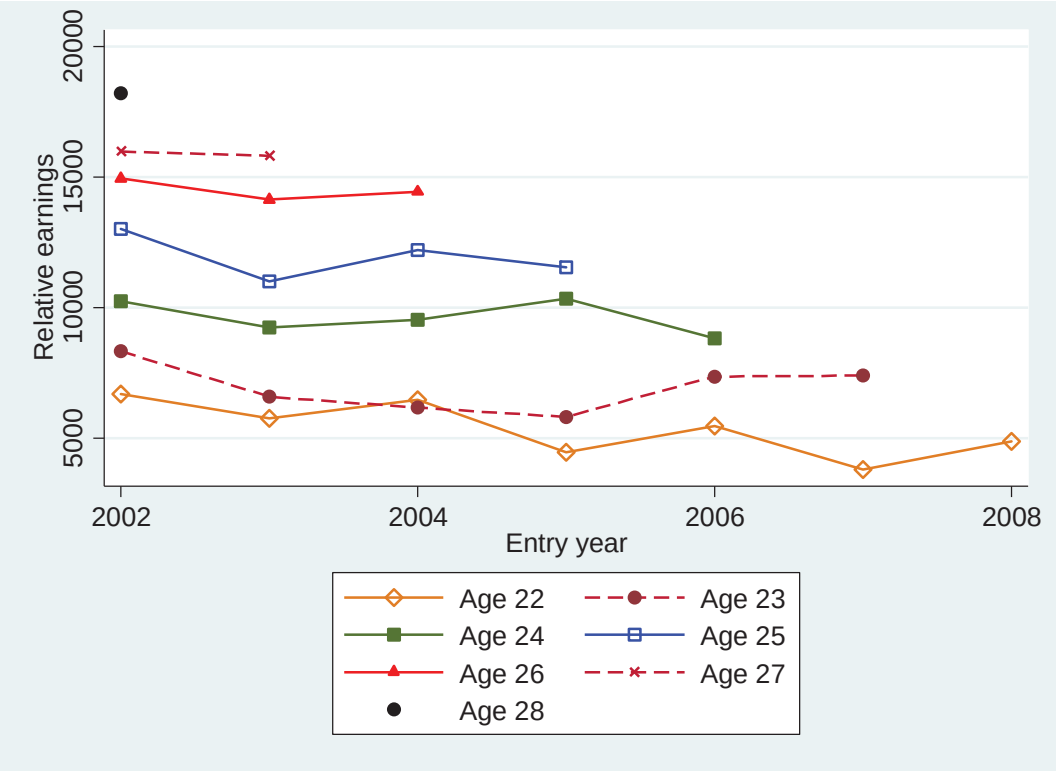

Figure 3 shows the results of this exercise. For every age, the later the entry year, the lower was the relative wage. The effect was greatest for entrants to the labour market: entrants in 2008 received 73% of the relative wage received by entrants in 2002, whereas those with one year of labour market experience received 89% of the relative wage if they had entered in 2007 instead of 2002. It is true that for each entry cohort, the wage rose with age (representing labour market experience) relative to the sample average wage of that year. Nevertheless, we see that new graduates suffered a decline in their wages relative to urban wages as a whole. New graduate wages have indeed been flexible in relative terms.

8.2 The effect on unemployment

Unemployment is defined as being in one of three of the 15 categories of employment status shown in the NBS questionnaire: job seeker, awaiting job allocation by the state, and other unemployed person. The employed comprise the seven categories of

employee and self-employed. The employed plus unemployed give labour force participants, and the unemployment rate is unemployment divided by labour force.

Figure 4 shows that graduate unemployment falls with age. For instance, the unemployment rate of a 22-year-old in 2002 was as high as 29%. With the exception of a rise in 2004, the rate fell monotonically for these 2002 entrants, being 18% at age 24 and 9% at age 26. For the 2003 entrants, the unemployment rate fell from 19% at age 22 to 12% at age 25. Despite the initial difficulty of finding jobs, it appears that the process of job search eventually provided employment for an increasing proportion of each graduate cohort as it aged.

Unemployment was generally most severe for 22-year-olds, i.e. those entering the labour force. The unemployment rates of 22 year-olds and 23 year-olds rose between the entry years 2002 and 2004. However, the unemployment problem for graduate entrants to the labour force eased after 2004. It is not the case that the ever-expanding number of graduate labour market entrants produced an ever-increasing number of new graduates without jobs.

We need not show a graph of the unemployment rate of graduates relative to that of all employees. As the latter stayed almost constant over the period 2002-2006, the absolute unemployment rate is also the relative unemployment rate.

Figure 4. Higher education graduates: unemployment by age for each entry cohort

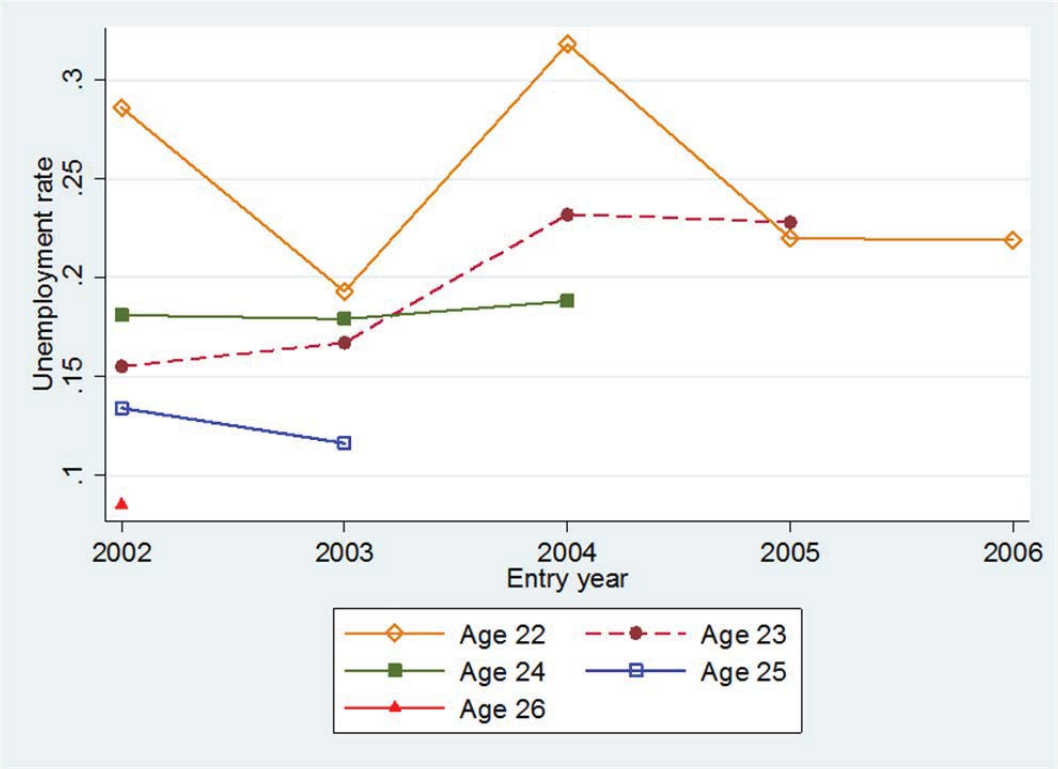

8.3 The effect on good jobs

The NBS questionnaire shows eight categories of occupation. Two of these – technical and professional, and government or enterprise leader - are by both description and average earnings most likely to represent 'good jobs'. We use that definition and express good jobs as a proportion of all categories of occupation.

Figure 5 depicts the proportion of new graduates in good jobs by entry year for each age. We see that the proportion of 22-year-olds in good jobs fell between the 2002 and 2003 entry, from 31% to 22%, but thereafter generally rose, reaching 28% for the 2006 entry cohort. Similarly, in the case of 23 year-olds the proportion rose with entry year for those who had entered after 2003. There is little sign that these young graduates filtered down into lesser jobs as their numbers expanded rapidly.

Figure 5. Higher education graduates: proportion in good jobs by age for each entry cohort

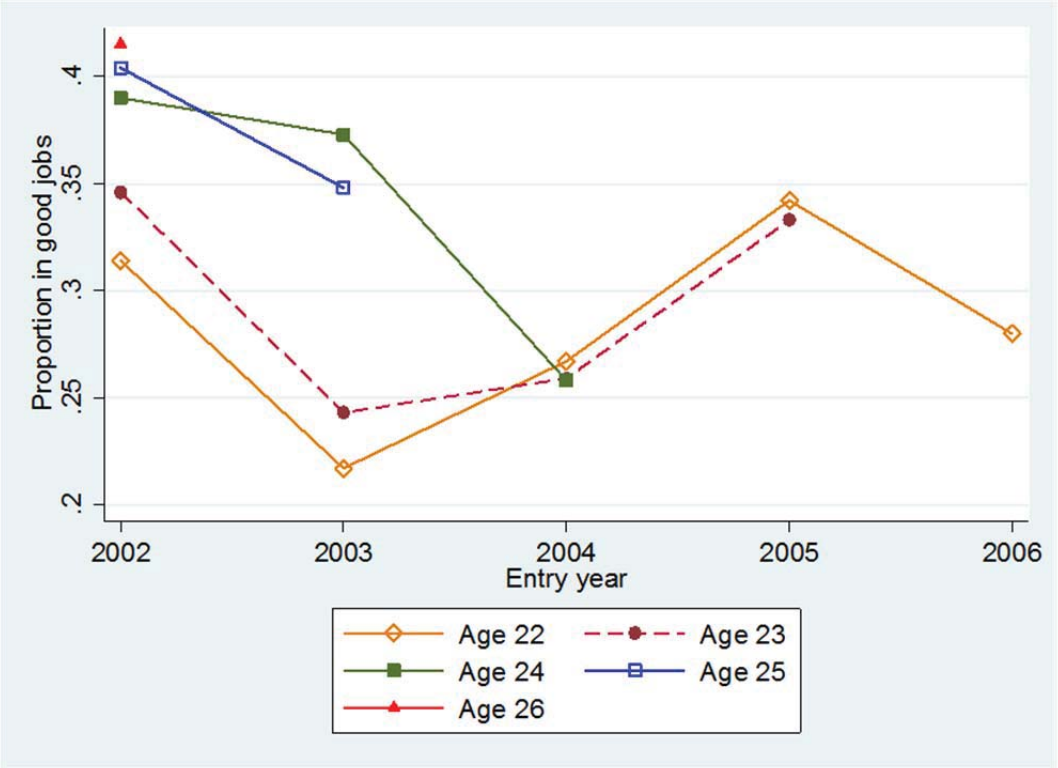

For the 2002 entry cohort, the proportion in good jobs rose monotonically, from 31% at age 22 to 39% at age 24 and to 42% at age 26. The pattern is similar for the 2003 entry cohort. This progression up the job ladder is to be expected, as workers gain experience and acquire human capital on the job. The great expansion of higher education did not obstruct the normal job promotion that comes with experience and skill acquisition.

It is unnecessary to analyse the proportion of new graduates relative to a comparator group: the proportion of all surveyed workers in good jobs stayed almost constant between 2002 and 2006, so the absolute proportion is also the relative proportion.

Is there inconsistency between the CHIP and NBS results on occupational attainment? The CHIP surveys provided evidence that the proportion of the graduate five-year entry cohort in 'good jobs' fell by ten percentage points between 2002 and 2007, from42% to

32%. The NBS surveys indicated that a low proportion of graduates in their entry year had access to good jobs, but that access improved if they entered after 2004..

Nevertheless, there are discrepancies between the two sources which might lie in different sampling methods, different definitions of good job, or different definitions of an entry cohort. It appears that CHIP used a different definition of the category 'managerial and professional', and that the definition might have varied from one CHIP survey to the other. The proportion of employed workers overall was 33% in CHIP 2002 and 29% in CHIP 2007.However, even if this four percentage point fall in the overall proportion in good jobs is to be discounted, we are still left with a six percentage point fall in the proportion of the graduate entry cohort in good jobs.

9. Conclusion

The CHIP surveys of 2002 and 2007 provide vivid evidence of the huge expansion of higher education over the five years. The rapid growth of higher education graduate employment as a proportion of total employment suggests that the higher education wage premium should have been depressed by market forces. However, a conventional wage function analysis shows that the wage premium fell only very little. There are various possible reasons for this result. The explanation which we explore is that the effect of the supply shock is selective in its effects, concentrating on the most recent cohort or cohorts. The same hypothesis can be advanced to explain the very small rise in the graduate unemployment rate.

In the case of higher education (both university and college) graduates we found a fall in the average hourly wage relative to high school leavers over the five years, but only for the entry cohort. The same is true of university graduates alone but the fall in the ratio is much greater. There is the same pattern for the unemployment rate: it rises only for

the entry cohort, and more so for university graduates than for higher education graduates. The third hypothesis was that new entrants filtered down into lesser jobs. Indeed, the proportion of new entrants in 'good jobs' fell over the five-year period, especially in the case of university graduates. In this case the fall was extended, more weakly, to the next five-year cohort.

Using the NBS urban sample surveys, when we examined the behaviour of real wages of new graduates from higher education relative to that of a comparator group, we found that the relative wages of the graduates fell dramatically as the entry of these graduates into the labour market rocketed. Large parts of the Chinese labour market have been, and may well still be, characterised by market imperfections and wage rigidities (Knight and Song, 2005). However, the NBS evidence indicates that this is not the case in the labour market for young higher education graduates.

Our other two variables to be explained are understandable in the light of this wage flexibility. The unemployment rate of graduates entering in 2002 was over one third. For later cohorts, however, the unemployment rate of graduate entrants fell. Moreover, the unemployment rate of each entry cohort fell sharply as the cohort aged. It appears that the fall in real wages eased to problem of unemployment. The entrants adapted quickly to the new situation and were willing to take the available jobs at the lower wage levels. We saw no cascading of unemployment to which – it might be hypothesised – the great expansion each year of labour market entrants with higher education would give rise.

This process of absorption might suggest that young graduates accepted the need to enter into lesser jobs as part of the labour market adjustment. Indeed, this appeared to happen for a year or two. The proportion of youngest graduates in 'good jobs' (as we

define them) was low throughout the period under review but rose with entry year beyond 2004. The relative fall in the cost of employing them could have helped to increase the demand for young graduates, and this might have raised the proportion in good jobs. The earliest entry cohorts increasingly entered good jobs as they aged, as would be expected. We found no evidence of a powerful process of continuing occupational downgrading which – it might be hypothesised - the great expansion of higher education would precipitate.

The objective of Section 7 was to investigate the extent to which the labour market for young graduates had been affected by the change in quality of individuals and institutions that was likely to have accompanied rapid expansion in enrolments. Our attempt to examine the effect of the expansion of higher education on quality had limited success owing to the lack of relevant data and the differences in the information available in the two CHIP surveys. Only one question on quality was present in both surveys: the self-reported quintile of academic performance in high school. It might be expected that employers had to discriminate among applicants of more varied quality after educational enrolments had expanded. We found no evidence for new entrants that good performance in high school was more rewarded after the expansion of enrolment than before, although this was true for graduates as a whole. The main evidence consistent with the hypothesis is that employers rewarded the high school performance of new entrants more than that of established workers and that the *gaokao* score (available for 2007) raised wages more for new entrants than for established workers.

We were unable to test the plausible hypothesis that the expansion of higher education had increased differences in labour market performance between high-quality and low-quality graduates, on the one hand, and graduates from high-quality and from low-

quality institutions, on the other. The conclusion from the analysis of the two CHIP surveys is that the direct effect of post-school expansion on graduate wages and its indirect effect through a consequent decline in quality could not be distinguished. Nevertheless, the analysis serves as a guide to the data and methods that would be required to make that distinction in future research.

General household or labour force surveys – such as those that we have analysed – are not ideal instruments for the research questions posed in this paper. A promising line of research would be to conduct 'tracer surveys' of graduates, using samples that are representative of the variety of academic institutions, disciplines, and years of entry to the labour force.

References

Appleton, Simon, John Knight, Lina Song and Xia Qingqi (2002). 'Labor retrenchment in China: determinants and consequences', *China Economic Review*, 13: 252-75.

Appleton, Simon, John Knight, Lina Song and Xia Qingxi (2004). 'Contrasting paradigms: segmentation and competitiveness in the formation of the Chinese labour market', *Journal of Chinese Economic and Business Studies*, 2, 3: 185-205.

Chi, Wei, Richard B. Freeman, and Hongbin Li (2012). 'Adjusting to really big changes in the labor market in China, 1989-2009', National Bureau of Economic Research, Working Paper 17721, January.

Freeman, Richard B. (1976). *The Overeducated American*, New York: Academic Press.

Gustafsson, Bjorn, Li Shi and Terry Sicular (eds) (2008). *Inequality and Public Policy in China*, New York: Cambridge University Press.

Knight, John (1979). 'Educational expansion, occupational production functions and filtering down', *Oxford Economic Papers,* 31, 2, July: 187-204.

Knight, John and Li Shi (2006). 'Unemployment duration and earnings of re-employed workers in urban China', *China Economic Review,* 17: 103-19.

Knight, John, Deng Quheng and Li Shi (2011). 'The puzzle of migrant labour shortage and rural labour surplus in China', *China Economic Review*, 22, 4: 585-600.

Knight, John and Richard Sabot (1990). *Education, Productivity and Inequality: the East African Natural Experiment*, Washington DC: World Bank and Oxford University Press.

Knight, John and Lina Song (2005). *Towards a Labour Market in China*, Oxford: Oxford University Press.

Knight, John and Lina Song (2008). 'China's emerging urban wage structure, 1995-2002' in Gustafsson, Li and Sicular (eds), *Public Policy and Inequality in China*, New York|: Cambridge University Press: 221-242.

Li Shi and Xing Chunbing (2010). 'China's higher education expansion and its labor market consequences', IZA Discussion Paper No. 4974, May.

Li Shi, Hiroshi Sato and Terry Sicular (eds) (2013). *Rising Inequality in China. Challenges to the Harmonious Society*, New York: Cambridge University Press.