

Immigration and the Canadian Earnings Distribution in the First Half of the 20th Century: Appendix

Alan G. Green*and David A. Green†

July, 2015

*Department of Economics, Queen's University (dec.)

†Vancouver School of Economics, University of British Columbia and International Research Fellow, IFS

1 Census Data Details

In this section of the Appendix, we provide further details on the Census data.

1.1 Occupations

The occupation variable is based on a single question in the 1921, 1931 and 1941 Censuses but in 1911 Census, the respondent was asked separately about their "chief" occupation and about other occupations. However, almost no workers reported an "other" occupation, leading us to believe that the data is directly comparable across years. In samples we drew directly from the 1911 manuscripts for another project we found that only 2 of 1895 observations in our random sample of males in Vancouver and 9 of 2393 observations in Montreal reported a secondary occupation. We also checked for differences by skill level by drawing all male respondents in Montreal who listed their occupation as machinists or labourers. None of the 2648 machinists and only 231 of the 25694 labourers reported a secondary occupation.

Occupation corresponds to jobs held at the time of the Census, while earnings and weeks worked correspond to all jobs in the previous twelve months. Thus, our calculated average earnings for an occupation may not be the true average earnings for people who work only in that occupation. For example, if workers employed in a skilled occupation at the Census date spent part of the year working as labourers then our calculated average earnings would be lower than the average earnings for a person employed solely in that occupation. Comparisons to other data indicate that this is not a concern for the least skilled workers (e.g., labourers) or most skilled (e.g., professors) in Census data but may be important when considering the earnings of tradesmen (Green and Green 2008).

1.2 Employment Definitions

The only substantive difference in variables of interest to us across the Censuses we use is in the way weeks of employment were obtained: in 1911 and 1941, respondents were asked directly about weeks of employment in the previous 12 months, while in 1921 and 1931 they were asked about weeks of unemployment, u , and employed weeks were calculated as $(52 - u)$. If respondents differentiated between weeks of unemployment and weeks in other non-employed states then calculated weeks of employment would tend to be over-stated in 1921 and 1931, and employed weeks do seem to be relatively high in those years. More importantly, while almost all male employees with positive reported annual earnings respond to the weeks of employment or unemployment questions in the 1911, 1931 and 1941 Censuses, approximately a third do not have a recorded weeks of unemployment in 1921. This leaves a somewhat difficult trade-off. We would prefer to focus on weekly wages, both because it is closer to the price of labour construct upon which we focus later in the paper and because it is more comparable to what has been used in earlier papers. However, the weekly wage measure is clearly inferior to the annual earnings measure. We respond by focusing mainly on earnings but providing estimates for weekly wages in this appendix for comparison. Also, in making comparisons across years, we use only the 1921-1931 and 1911-1941 pairs when examining weekly wages since these are the pairs of years with the same employment questions. As we will see, most of our conclusions are unchanged whether we use earnings or weekly wages.

¹

¹Because of rapid deflations in both 1921 and 1931, choosing to average in this way yields quite different results from just using one year's index value in each case. Thus, the actual index values we use for Montreal (compared to a 1913, Toronto base of 100) are 83 for 1911, 148 for 1921 and 127 for 1931. If, instead, we had used the values for 1911, 1921 and 1931, the index values would have been 82, 136 and 116, respectively, thus affecting the 1911-1921 comparisons. However, we believe that the averaging approach is the most reasonable given the timing of the earnings

2 Comparisons to other Canadian Wage Series

The two series that cover much of the same time period as our Census data are the hourly wage series in Emery and Levitt(2002), based on Labour Gazette data, and the wage data from the Canadian Pacific Railway records collected by Mary Mackinnon (Mackinnon(1996)). It is worth noting at the outset that there is some degree of controversy about the reliability of both data sources. As Mackinnon(1996) discusses, the Labour Gazette data for the metal trades, printing trades and building trades are likely union scales and may not have corresponded to what was paid to non-union workers, or even to what was actually paid to unionized workers.² However, Altman(1999) argues that some of the non-metal manufacturing data corresponds more to non-union workers so series such as Emery’s common factory labourer wages, constructed as average wages across labourers in manufacturing firms reporting to the Labour Gazette, may be more representative. On the other side, the CPR data comes from company pension-related records. As such, they are likely to be accurately recorded but might be questionable in terms of their representativeness relative to the rest of the workforce. Mackinnon(1996) examines and rejects the main potential objections to the data on these grounds but she also states that government control of railway wages and prices between 1917 and 1921 led to disproportionately large increases in wages in that period relative to the rest of the workforce. In fact, even in 1925, the Labour Gazette reports that wages in all the principal railways ”are fixed according to agreements between the several railways and the organization of railway employees.” (Canada (1925)) implying that it is actually the railway data that reflect union scales. Mackinnon herself states that, ”From 1918 on, important institutional factors affect railway wages, so that CPR data are a better indicator of general wage trends for the earlier period ...”(Mackinnon(1996, p. 115). Both Mackinnon(1996) and Emery and Levitt(2002) provide comparisons across the various data sources, including the Census. We will discuss their conclusions as we proceed.³

Since both the Gazette and CPR series correspond to hourly wages, we need to convert the Census data from weekly to hourly wages to make comparisons. Herb Emery has collected hours series from the Gazette. We have access to the data on hours for common factory labourers and machinists, allowing us to create comparisons for a standard type of skilled/unskilled wage ratio. In order to control for differences in regional coverage across the dataset, we focus our attention on Montreal.

Appendix Figure 6 contains a plot of Emery and Levitt(2002)’s common labourers in factories wage series for Montreal. ⁴ This series was constructed by collecting all the plant specific wages for each city reported in the supplements to the Wages and Hours publications. The factory labourers series actually ends in 1938. We extend it to 1940 using the trend from Emery’s building labourer’s series, which is virtually flat in 1938, 1939, and 1940. Note, also, that the years 1922-1924, 1931 and 1934 are not available in the Gazette data and, for the purposes of the figure, we use a simple linear interpolation for those years. We also plot the implied hourly wage rates for Montreal from the Census data using Emery’s hours per week for common factory labour to convert to hourly wages. We plot the Census numbers with a linear interpolation linking each point. In all cases, we deflate the wages using Emery and Levitt(2002)’s price index for Montreal and report all numbers

reporting. Since we apply the same index values to all earnings levels, using different index values alters the relative location of each year’s wage distribution but not our conclusions about inequality movements.

²Urquhart and Buckley(1965) report that unionized labour made up 8% of the non-agricultural workforce in 1911, 14% in 1920 and 11% in 1931.

³Another point of comparison is found in Meltz and Stager(1979), who examine wages in 52 occupations across the 1931,41,51,61 and 71 Censuses. They argue that the 1931-41 and 1941-51 periods are both characterized by compression. As we will see, this fits with the patterns from both the Census and Gazette data.

⁴We are grateful to Herb Emery for providing us with his common labourers in factories series.

in 1913 real Toronto dollars to match our reporting of the Census numbers above. The third line plots labourer wages from Mackinnon's CPR data for the period up to 1930 (since the CPR data does not extend to 1940).

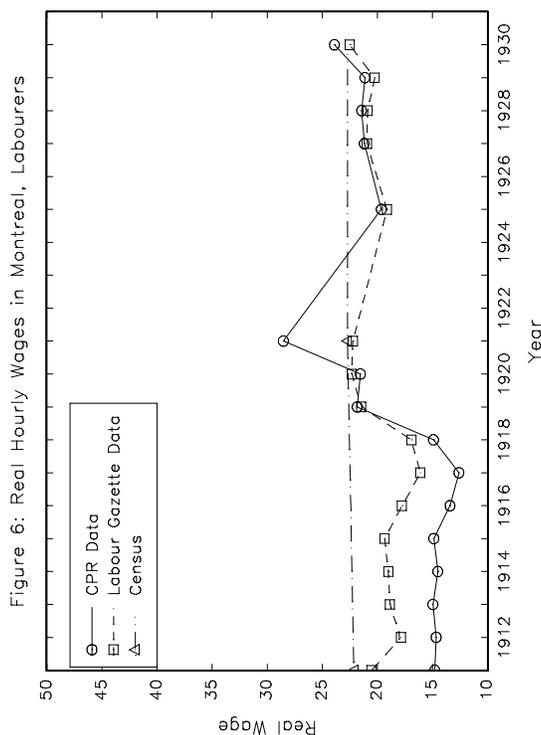


Figure 6: Real Hourly Wages in Montreal, Labourers

Two features stand out in this figure. The first is the relatively high value for 1921 in the CPR data. This occurs because the official nominal wage for labourers actually increases in a year with substantial deflation. The second feature is the relative values of wages from the three sources. Putting aside the seemingly anomalous result from the CPR in 1921, which may be a reflection of special contracting conditions at the railway, the interpolated Census value and the observed values for the Labour Gazette and CPR data are extremely close to one another in 1920 and in 1930. The Census and Labour Gazette data are also very close in 1911, with both being substantially above the CPR wage. This pattern echoes a remark made by Emery and Levitt(2002). Commenting on Mackinnon(1996)'s claim that weekly wages in the 1911 Census may have been abnormally high because workers were putting in overtime, Emery and Levitt note that dividing Census weekly wages by their common factory labour hourly wage yields implied hours per week that are very close to those reported in the Labour Gazette.

There are some other sources against which to compare the various wage series. Mackinnon(1996) reports daily wages for labourers employed by the government to work on the canals around Montreal. In 1911, these reported wages were \$1.50 per day. If these labourers worked 9 hour days then this would correspond to an hourly wage (in Toronto dollars) of approximately 17

cents per hour. This is between the CPR (15 cents) and the Census and Gazette numbers (22 and 21 cents, respectively). The Labour Gazette in February 1912 reports that 1,800 street labourers working for the city of Montreal received an increase in pay from \$2.00 to \$2.10 per day in January, 1912 (Department of Labour(1912)). If we then assume that the \$2.00 figure is relevant for 1911 and again assume a 9 hour day, the implied hourly wage (in 1910/11 Toronto dollars) is 0.23. If we assume they worked a 10 hour day then the relevant hourly wage is 0.21. In either case, the hourly wage is very close to those derived from the Census data and reported in the common factory labourers series from the Labour Gazette.

In the end, the congruence of two such different sources (the Gazette data and the Census) combined with the evidence on street labourers' wages and the evidence that the CPR data seems to be habitually below other sources in the pre-war period suggests to us that we should put more credence in the Gazette and Census wages. Those sources indicate an essentially flat real wage for labourers between 1910 and 1930.

In Appendix Figure 7, we repeat this exercise but examine machinists wages. Machinists are a skilled trade for which wages are readily available from all three sources.⁵ Both Emery and Levitt(2002) and Mackinnon(1996) emphasize the broad agreement of various sources on skilled wages. All three sources are in close agreement on the real wage for this occupation in 1911 and all three suggest relatively substantial increases in machinists real wages between 1911 and 1930/31. However, there are also some strong differences among the series. In particular, both the CPR and Gazette data show large increases in machinists real wages between the pre-war period and the immediate post-WWI period. In contrast, the Census data indicates that the real wage for machinists was essentially unchanged between 1911 and 1921. It is worth noting that the Labour Gazette data is again closer to the implied Census data than the CPR data. The differences between the implied Census hourly wages and the other two series in 1921 might be accounted for by the extremely tumultuous nature of the labour market in the immediate post-war years. Hours per week changed dramatically in the span of a few years in this period. There was also rapid deflation in 1921. The common factory labourers' wage series tracks this deflation quite closely while the more skilled workers' official wages follow with a lag. It is possible that the relative lack of flexibility in machinists' wages induced other adjustments, such as in hours of overtime available. In the longer run, though, the Census data also shows increases in the real wage of machinists to a degree comparable though not as large as what is observed in the Labour Gazette data.

One reason to examine the sources other than the Census is that this allows us to compare wages at cyclically similar points. The 1911 Census was taken at the time of a boom while the 1921 and 1931 Censuses correspond to a recession and a depression of differing severity. To get more cyclically constant results, we could compare the real wages in 1911 with a prosperous year from the 1920s such as 1927. The common factory wage labourers real wages in those two years are very similar (\$0.20 and \$0.21, respectively). For machinists, the Labour Gazette data shows a marked increase from \$0.31 to \$0.38. Thus, the picture from the Census that skilled wages had increased substantially while unskilled wages had changed little from the pre-war period to the end of the 1920s appears to hold up to comparison at cyclically similar points.

⁵The Labour Gazette publications on hours and wages actually include two different versions of the machinists' wages for Montreal for 1921. In the reports near the actual date (Reports number 4 and 6), the wage range for machinists is listed as, .55 to .70 cents per hour. However, starting with report 7 the range is listed as .55 to .90. We used the mid-point of the former range because this is in closer accord with other evidence, particularly machinists' wages in Toronto. Throughout the Labour Gazette data the Toronto wage ranges have mid-points very similar to those for Montreal. The one exception to this is the 1921 range listed for Montreal in Reports 7 and later. In 1921 the listed wage range for machinists in Toronto is .50 to .75. Thus, we believe that data was added later on a plant that was an outlier and stick with the earlier listed data.

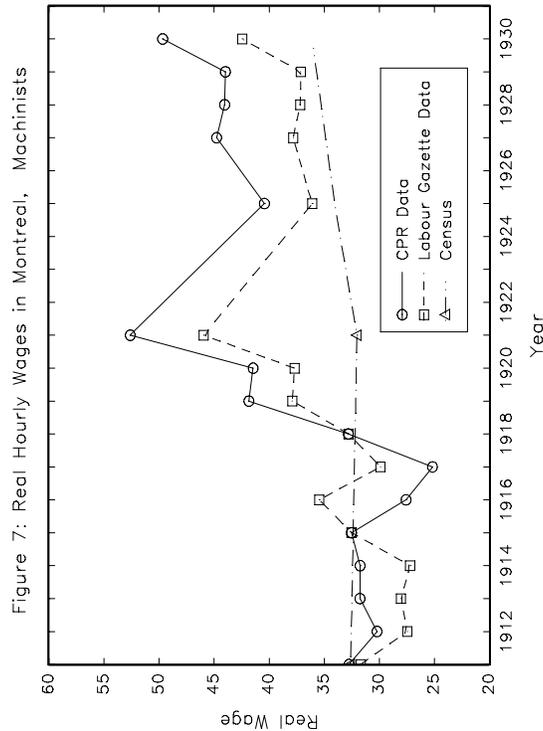


Figure 7: Real Hourly Wages in Montreal, Machinists

Working with the labourer and machinist series just described, we can generate a skilled to unskilled hourly wage ratio that is similar to what is presented in earlier papers. These ratios are presented in Table 1.

The three series show somewhat different patterns. Both the Census and the CPR data show a drop in the skill ratio between 1910 and 1920 followed by a strong increase between 1920 and 1930. However, the initial drop is much larger in the CPR data. As we discussed earlier, the labourer wages in the CPR data for 1921 appear to be out of line with other evidence and so it is unlikely that this data reflects general trends. For the Gazette data, we see the same long run pattern of increase in the skill ratio as is observed in the Census data but with very different timing across decades. In our view, the most reliable trends come from the Census. The wages in the CPR often appear to be out of line with other sources and, at times, reflect contracts and government mandates rather than general wage conditions. The Gazette appears the over-represent union agreements in an era when union membership was did not make up a large portion of the workforce. In this regard, it is reassuring that for labourers' wages (where unionisation was likely not as much of an issue) the Gazette and Census are in strong agreement. We also gathered information on Assistant Professor Salaries for the University of Toronto (averaged across all departments) from the Ontario Sessional papers. Taking the ratio of those salaries to 52 times the average weekly wage for labourers in Toronto from the Census, we get earnings ratios of: 2.63 in 1921, 3.81 in 1931, and 3.0 in 1941 (fitting with the general pattern of strong increases in inequality in the 1920s followed by smaller

Table 1: Machinist to Labourer Hourly Wage Ratios

Year	Census	Gazette	CPR
1910	1.46	1.54	2.18
1920	1.42	1.87	1.88
1930	1.62	1.88	2.08
1940	1.6	1.78	-

Sources: Census data is restricted to Montreal. Conversion to hourly wages is done using Emery's Labour Gazette hours series. The Labour Gazette data is a continuous series underlying the data reported in Emery and Levitt(2002). Gazette data in the 1940 row is actually for 1938. Common factory labourers wages are the denominator for both Gazette series. Census building trades are an average of all building trades while the Gazette data corresponds only to carpenters. The CPR series is from Mackinnon(1996) and is based on Canadian Pacific Railway administrative series.

declines in the 1930s). In comparison, using average earnings for professors in Toronto from the corresponding Censuses we get ratios to labourer earnings of: 2.60 in 1921; 3.95 in 1931; and 2.81 in 1941. Thus, the Census is in strong accord with another data source for non-union workers at the top of the earnings distribution. Our overall conclusion is that the Census series differ from the other two available wage sources but that there are good reasons to prefer the Census numbers.

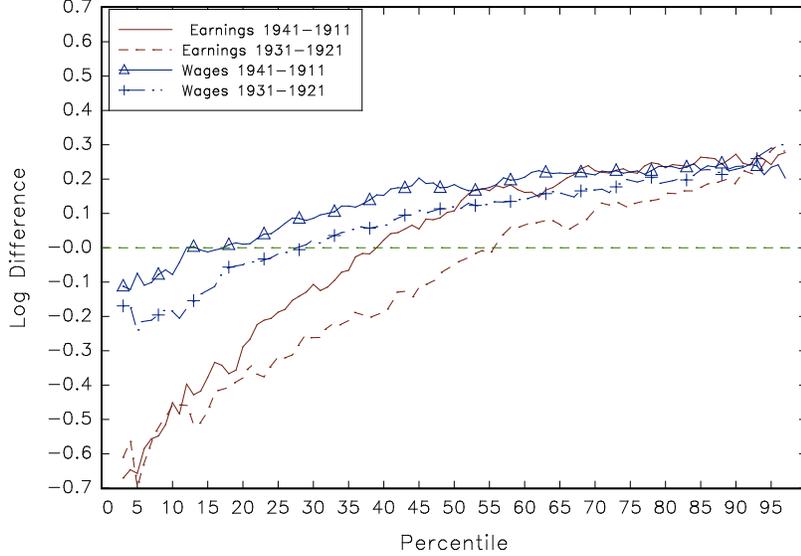
3 Weekly Wages and Annual Earnings

In Figure 4, we show the 1911 - 1941 difference for annual earnings and reproduce the 1921-31 difference. We focus on these differences in order to compare to the weekly wage differences that are presented in the same figure. Recall that these two differences are the ones that are the most reliable for weekly wages because of changes in questions about weeks of work over time. The 1911-41 difference for earnings shows that there is a strong long-term increase in inequality below the median but a relatively constant increase in earnings levels across percentiles above the median. For weekly wages, the long term pattern has a roughly similar shape to annual earnings, though much more muted and with increases in level across much of the distribution. Whereas with annual earnings there are declines at virtually all percentiles below the median, in weekly wages there are declines only below the 15th percentile. As with annual earnings, the increase in inequality in the weekly wage is driven mainly by movements in the 1920s.

One immediate point of interest is the much more substantial decline in the lower half of the annual earnings distribution than is evident in the weekly wage distribution. Given that weekly wages are constructed as annual earnings divided by weeks of work, the difference must reflect a drop in weeks of work. Indeed average weeks worked per year declined from 44 in 1921 to 42 in 1931. This may reflect selection issues related to the fact that approximately a third of paid employees in the 1921 Census do not report their weeks of work.⁶ However, this would not explain the substantial difference between wage and earnings patterns for the 1911-1941 change since there are no such selection issues in either year. In the end, given the changes in weeks worked questions and

⁶A decline in weeks of work from 1921 to 1931 may also seem surprising given the labour market problems in 1921. However, while 1921 was a bad labour market year, the labour market in the summer and fall of 1920 was actually strong (exactly when it began to turn bad - either in October or in December - depends on whether the source of information is unions or employers (Dominion Bureau of Statistics (1921))). The 1921 Census asks about labour market experiences over the 12 months preceding the start of June Census date. Thus, 7 out of the 12 months are from 1920 and, as a result, the 1920/21 Census data does not represent as bad a labour market as one might expect given the 1921 Census date. The 1930/31 Census data represents a similarly mixed year.

Figure 4: Differences in Log Percentiles,
Male Earnings and Weekly Wages



the possible selection effects in 1921, we conclude that it is preferable to focus on annual earnings, where we have a consistent measure over time.

4 Reduced Form Wage and Earnings Change Regressions

In this section we present the results for both weekly wages and annual earnings of a reduced form approach to assessing the impact of immigration on the national wage and earnings structure. This approach consists of regressing the change in the log of native born earnings on a variable representing an immigration "shock" in skill cells defined by occupation and age. That is, with occupation indexed by j and age indexed by k :

$$\Delta \ln(w_{Njkt}) = \alpha_{0t} + \alpha_1 \frac{\Delta E_{Ijkt}}{E_{Njkt-1} + E_{Ijkt-1}} + \Delta u_{jkt} \quad (1)$$

where N refers to native born workers and I refers to immigrant workers, the E 's are employment levels, u_{jkt} is an error term, and α_{0t} corresponds to year effects. Borjas(2003) derives an equation similar to this one assuming a linear demand function for workers of a type defined by the j and k dimensions along with a linear supply equation for native born workers. Notably, native born and immigrant workers are assumed to be perfect substitutes and each skill cell is treated as a separate market in the derivation. The properties of the error term, u , are crucial in determining the estimation approach to equation (1). Borjas(2003)'s approach (again, translated into our context) allows for time invariant occupation, age and age-occupation interaction effects and also for general time effects and time trends, separately, in the occupation and age effects. Estimating in differences within cells, this amounts to having $\delta u_{jkt} = \phi_k + \phi_j + e_{jkt}$, where ϕ_k and ϕ_j are occupation and

age group specific effects, respectively, and e_{jkt} is a white noise error term. The key identifying assumption is that once one controls for age and occupation effects in this way, the remaining variation in the error term is uncorrelated with changes in the ratio of immigrants to the native born in the cell. That is, immigration is viewed as exogenous.

While it is common to approach this estimation treating immigration as an exogenous supply shock within a given cell, it is also common to be concerned about potential remaining endogeneity. In particular, in our context, we know that immigrants who were explicitly brought in to work in agriculture in the West, were, by the time we observe them in the next Census, in a variety of occupations and locations. Thus, the set of changes in immigrant employment we observe might reasonably be expected to be partly due to endogenous responses to within-cell productivity shocks. One approach taken to this problem by a number of authors, starting with Card(2001), is to instrument for the change in immigrant supplies using a variable based on the notion that immigrants will move to locations within the host country where there are enclaves of people from their country of origin, regardless of the local economic circumstances. We use a variant on this instrument in which we assume that immigrants move to a geographical location based partly on these enclave considerations and then, once there, pursue the most common occupations in that location. That is, we assume that when they search for a job in the location they are attracted to for companionship reasons, their probability of finding a job in a given occupation is proportional to the size of that occupation in that location. Moving to a mining town, immigrants are more likely to find mining jobs than if they moved to Toronto, say.

To construct the instrument consider an immigrant from source country g , and let the proportion of immigrants from g in location r in year t be given by p_{gtr} . Also, let the proportion of all workers in occupation j in location r in period t be given by π_{krt} . We will predict the number of people from country g in occupation k in year $t+1$ at the national level as:

$$n_{kgt} = \sum_r \pi_{krt} * p_{gtr} * \Delta n_{gt} \quad (2)$$

where Δn_{gt} is the change in the number of immigrants from country g at the national level between t and $t+1$. Note that this exercise is done separately by age group with both the π_{krt} weights and the national level inflows (Δn_{gt}) being age group specific.⁷ The key identifying assumption is that neither the composition of employment by occupation at the start of a period in a location nor the distribution of immigrants from a given country across locations at the start of the period are correlated with the change in productivity within an occupation-age cell at the national level in the ensuing decade. We also require that the general growth in immigration from each country is not correlated with the cell specific productivity shocks. These assumptions would be violated if people from a given source country only immigrated to Canada because there was an increase in demand for a specific occupation in a specific location. The fact that π_{krt} is calculated for all workers in the location, not just for immigrants from the specific source country, makes it more likely that the requirements for the instrument to be valid are met.

We estimate the specification given in (1) at the age by occupation group level for the 1911-21, 1921-31 and 1931-41 differences pooled using real annual earnings. We employ two different

⁷The regions are defined as follows. We divide Canada into rural areas, towns, and 13 major cities (Victoria, Vancouver, Edmonton, Calgary, Regina, Winnipeg, Hamilton, Toronto, Ottawa, Montreal, Quebec City, St. John, and Halifax). Towns are all locations with populations over 5000 that are not on the list of major cities, and rural areas are all remaining locations. We define a region as either a major city or a town or rural area within a given province (e.g., towns and rural areas in British Columbia are two of our regional groups). We group Prince Edward Island and Nova Scotia together, resulting in 8 provincial groups. Thus, the total number of regions is $13 + (8*2) = 29$.

occupation categorizations. In the first, which we will call the narrow grouping, we calculate the number of workers and the average wage in cells defined by the 4 age groups and the 146 occupation groups described earlier. In the second, which we will call the broad grouping, we work with cells defined by the 4 age groups and 10 broader occupation groups.⁸ The narrow occupations allow for more observed wage and employment ratios but those observations are also noisier than with the broad occupations. We drop cells with fewer than 10 observations and weight using the square root of the cell size. The dropping of small or empty cells leaves 1534 observations in the narrow occupation specification and 120 observations for the broad occupation specification. In all of our estimations, standard errors are clustered at the occupation x age group level to allow for flexible forms of serial correlation in the error term. The regression includes an indicator for each period and a set of occupation dummies and age dummies.

We present estimates of α_1 in Table 3. We employ two versions of our instrument. The first (IV1) bases π_{krt} on the occupational distribution of all workers in region r while IV2 bases it only on the distribution of native born workers. We view the second as intuitively more likely to be valid but it is also more often weak.⁹ In the leftmost set of four columns, we work with the change in log real annual earnings while the rightmost four columns uses the change in log weekly earnings as the dependent variable. Using annual earnings, the basic OLS estimates of α_1 are -.061 and -.32 for the narrow and broad occupation groupings respectively, with the coefficient being statistically significant for the broad group estimation but not with narrow groups. A comparison with the estimates in the fifth column indicate that there is little difference between using annual earnings or weekly wages as our dependent variable, implying that the immigration shock had no impact on weeks worked. To the extent we believe the weekly wage numbers this would mean that immigration does not explain any of the large difference in the extent of declines in annual earnings relative to weekly wages. In the second and third columns, we present results using IV1 and IV2 in instrumental variables estimation. For the narrow occupation grouping, the first stage F-statistics associated with the instruments are just below the common weak instrument benchmark of 10. The resulting estimates of α_1 are much larger than the OLS in absolute value but are also much more poorly defined. For the broad occupation groups, the first stage F-statistics are over 20. Interestingly, the instrument based just on the native born occupational distributions at the regional level is, if anything, stronger. The estimates of α_1 using these instruments are both closer to, and not statistically significantly different from zero. Our conclusion from this table is that there is evidence of a basic negative correlation between immigrant changes in employment and native born wages but instrumental variables estimates from the broad groupings (where there is no weak instrument problem) indicate this effect is very close to zero.

One weakness of the approach to this point is that it treats changes in immigrant employment in an occupation identically whether the change results from movements of new immigrants or previous arrivals. But, as we have seen, there is good reason to believe that more established immigrants behave quite similarly to the native born in this era, suggesting that it is less plausible to treat their occupational changes as an exogenous supply shock. In the fourth and eighth columns of the table we break the immigrant change variable into changes due to immigrants who arrived in the decade corresponding to the change in the dependent variable (i.e., immigrants arriving between 1921 and 1931 for the observations with the 1921 to 1931 change in log earnings or wages) and immigrants who arrived in all previous decades. In all specifications, the effect of immigrants who arrived during the current decade is statistically significant and approximately three times the estimated effect with all immigrants combined. In contrast, the movements of earlier arrivals are

⁸The broad occupation groups are simple aggregations of the narrow groups and consist of: farm labourers, farm managers, general labourers, service, sales, semi-skilled, trades, clerical, managers and officials, and professionals.

⁹The precise version of IV1 is actually n_{kgt} divided by $E_{Njkt-1} + E_{Ijkt-1}$, and IV2 is formed analogously.

Table 2: Reduced Form Regression Results

	Narrow Groups				Groups			
	Earnings				Wages			
	OLS	IV1	IV2	OLS	OLS	IV1	IV2	OLS
Change in Immigration	-.061 (.043)	-.87* (.30)	-.48 (.32)		-.070 (.046)	-.64* (.24)	-.37 (.30)	
Immigration Last 10 Yrs				-.35** (.10)				-.22** (.056)
Immigration Before Last 10 Yrs				.12 (.051)*				.045 (.040)
No. of Obs	1534	1534	1534	1020	1507	1507	1507	993
First Stage F		9.9	9.2			10.05	9.4	

[htb]

	Broad Groups				Groups			
	Earnings				Wages			
	OLS	IV1	IV2	OLS	OLS	IV1	IV2	OLS
Change in Immigration	-.32* (.14)	-.11 (.16)	.16 (.16)		-.36* (.046)	-.027 (.20)	.22 (.21)	
Immigration Last 10 Yrs				-1.0** (.21)				-.81** (.19)
Immigration Before Last 10 Yrs				.30 (.19)*				.20 (.18)
No. of Obs	120	120	120	80	120	120	120	80
First Stage F		20.8	28.1			20.8	28.1	

Sources: The dependent variable is the change in log wages for native born workers in an occupation by age cell. The top panel is based on the narrow occupation grouping (146 occupations) and the lower panel is based on the broad occupation grouping (10 occupations). All specifications include a complete set of time, age, and occupation dummy variables.

Standard errors in parentheses. Standard errors are clustered at the occupation x age group level.

*, ** statistically significant at the 5, 1% level.

positively related to wage changes, suggesting that the earlier arrivals moved toward occupations with increasing demand.¹⁰

How large are these estimated effects? The broad group OLS estimates using all immigrants combined is most similar to the specification used in Borjas(2003)¹¹ and the estimate of -.32 is very close to his estimates using more recent US data. However, the estimated effect from recent immigrants is much larger. For the farm labour occupation, which we have seen was hit with a substantial immigration shock in the 1920s, the immigration shock variable on the right hand side of 1 takes a value of .19, which translates directly into a .19 log point decline in the wages for native born farm labour. Compared to an actual decline of .3 log points, this suggests that immigration could have played a substantial role in determining the wage outcomes in this broad occupation group. However, Ottaviano and Peri(2010) argue that α_1 has a very restricted (and possibly uninteresting) interpretation within a broader model of labour demand that incorporates imperfect substitutability among skills. We view the estimates in this section as useful in summarizing whether there appears to be a basic negative correlation in the data between immigration and native born wages. That does appear to exist, though the results are more convincing in this regard for the broad groups.

5 Complete Presentation of Wage Equation Estimates

5.1 Nested CES Model

We address the impact of immigration on the wage structure using a quasi-structural estimation approach based on a nested Constant Elasticity of Substitution (CES) production function. Our specification follows closely those in Ottaviano and Peri(2010) (hereafter, OP) and D’Amuri et al(2010) (hereafter, DOP). The advantage of this approach is that it allows for a complete accounting of the impacts of immigration by incorporating the possibility that an increase in the supply of immigrants in one occupation can have impacts on wages for other workers through substitutabilities and complementarities in production. Of course, the CES specification, even in its nested form, imposes significant restrictions on allowable patterns of interactions of use of factors. However, more flexible alternatives are not feasible in our situation with a limited number of years of data and a large number of skill groups.

We begin by considering the aggregate production specified as a simple Cobb-Douglas function of a labour aggregate, L_t , and capital, K_t :

$$Y_t = A_t L_t^\alpha K_t^{(1-\alpha)} \quad (3)$$

where Y_t is aggregate output and A_t is a productivity shifter.¹² We will assume that the economy faces fixed world prices for the output but that production also involves scarce factors (land in the

¹⁰In attempts to estimate the latter specification with instrumental variables, it was difficult to come up with a counterpart to n_{kgt} for the earlier arrivals and any instruments we tried had substantial weak instruments problems. Given that, we do not report IV estimates here.

¹¹Borjas(2003) uses 4 broad education groups and 8 age groups, implying similar relative sizes in cells to those in our broad group specification.

¹²In an earlier version of the paper, we allowed, further, for separate production functions in rural, major city, and town sectors. We have also estimated separate specifications in which we dropped the Maritime Provinces (which were recipients of small numbers of immigrants) and for the West alone. The estimated elasticities of substitution were very similar by sector and region, and so we have chosen, instead, to simplify to one aggregate production function.

rural sector, and entrepreneurial ability in the other sectors), implying downward sloping labour demand curves.¹³

Following DOP, we specify the labour aggregate as a CES function of types of skilled (or unskilled) labour. In their model, these skills are defined by education level while in ours they are the 10 broad occupations used in the investigations in the previous sections. Given this specification, the labour aggregator in period t can be written as:

$$L_t = \left(\sum_{k=1}^{10} \theta_{kt} L_{kt}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad (4)$$

where L_{kt} is employment in occupation k , and θ_{kt} are skill-time specific productivity shocks. The parameter σ determines the degree of substitutability among broad occupation groups.

We will express the labour in each broad occupation group as a CES aggregate of employment in each of the narrow occupations within the broad group (e.g., for different specific sales occupations within the broad Sales group), that is,

$$L_{kt} = \left(\sum_{l=1}^{n_k} \theta_{klt} L_{klt}^{\frac{\sigma_o-1}{\sigma_o}} \right)^{\frac{\sigma_o}{\sigma_o-1}} \quad (5)$$

where, n_k is the number of narrow occupations in broad group k , θ_{klt} are narrow occupation specific productivity shocks, and σ_o is the elasticity of substitution among narrow occupations within the same group. Working with the occupational employment in two levels in this way allows for the possibility that, for example, two types of sales occupations are more substitutable than sales workers and trades workers.

The skill specific labour can itself be written as a CES aggregate of employment for each of the 4 possible age groups described earlier:

$$L_{klt} = \left(\sum_{j=1}^4 \theta_{kljt} L_{kljt}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}} \quad (6)$$

where L_{kljt} is employment in narrow occupation l within broad occupation group k , age group j , and θ_{kljt} is the associated productivity shock. Again, ρ is a parameter capturing substitutability; in this case among age groups within the same narrow occupation.

Finally, within each type of labour defined by occupation and age group, we can observe workers broken down into new immigrants (who arrived in Canada in the decade preceding a given Census year, t), old immigrants (who arrived over ten years before t), and the native born. As we saw in section 2, Canada experienced inflows in immigration before 1921 that were both very large relative to its existing labour force and had an occupational distribution that was similar to that of the native born. The latter similarity increased with time such that a decade after arrival, the immigrant occupational distribution was very close to that of the native born in the three Censuses where we observe year of arrival (1911, 1921, and 1931). In reduced form estimates not presented here, we also observe that it is new immigrants who have the main impact on the wages of the native born. Based on this, we implement a specification in which we assume that old immigrants and the native born are close substitutes while new immigrants are a potentially more dissimilar factor. In terms of the nested CES specification, this implies two further levels of nesting: one at which a combination of old immigrants and the native born are treated as a separate factor relative

¹³This is in contrast to the specification for the manufacturing sector in Chambers and Gordon(1966) and makes our set-up closer to that in Lewis(1976).

to new immigrants, and another which is concerned with the aggregation of old immigrants and the native born. Thus, we will first write,

$$L_{kljt} = (\theta_{I_{kljt}} I_{kljt}^{\frac{\delta-1}{\delta}} + \theta_{ON_{kljt}} ON_{kljt}^{\frac{\delta-1}{\delta}})^{\frac{\delta}{\delta-1}} \quad (7)$$

where, I_{kljt} corresponds to employment of new immigrant workers, ON_{kljt} corresponds to employment of a combination of old immigrants and the native born, and δ is a parameter governing substitutability between new immigrants and old immigrant/native born workers in the same occupation-age group.

At the final level of nesting we have,

$$ON_{kljt} = (\theta_{O_{kljt}} O_{kljt}^{\frac{\phi-1}{\phi}} + \theta_{N_{kljt}} N_{kljt}^{\frac{\phi-1}{\phi}})^{\frac{\phi}{\phi-1}} \quad (8)$$

where, O_{kljt} represents employment of old immigrant workers, N_{kljt} represents employment of the native born, and ϕ is, again, the relevant substitutability parameter.

Using equations (3) through (8), and again following DOP closely, we can derive a wage equation for a native born worker in a given occupation-age cell, assuming their wage equals their marginal product:

$$\begin{aligned} \ln w_{N_{kljt}} = & \ln\left(\frac{\partial F}{\partial L_t}(K_t, L_t)\right) + \frac{1}{\sigma} \ln L_t + \ln \theta_{kt} - \left(\frac{1}{\sigma} - \frac{1}{\sigma_o}\right) \ln L_{kt} + \ln \theta_{klt} - \left(\frac{1}{\sigma_o} - \frac{1}{\rho}\right) \ln L_{klt} \\ & + \ln \theta_{kljt} - \left(\frac{1}{\rho} - \frac{1}{\delta}\right) \ln L_{kljt} + \ln \theta_{ON_{kljt}} - \left(\frac{1}{\delta} - \frac{1}{\phi}\right) \ln ON_{kljt} + \ln \theta_{N_{kljt}} - \frac{1}{\phi} \ln N_{kljt} \end{aligned} \quad (9)$$

We can similarly derive the log wage equations for an old immigrant worker in the same cell.

5.1.1 Estimating Equations and Results

The parameters in the nested CES model can be obtained sequentially, moving up the levels of aggregation. Thus, following OP, we can obtain an estimate of ϕ by taking the difference between the log wage expressions for native born and previously arrived immigrant workers:

$$\ln\left(\frac{w_{N_{kljt}}}{w_{O_{kljt}}}\right) = \ln\left(\frac{\theta_{N_{kljt}}}{\theta_{O_{kljt}}}\right) - \frac{1}{\phi} \ln\left(\frac{N_{kljt}}{O_{kljt}}\right) \quad (10)$$

This yields a simple intuitive formulation: holding constant relative differences in productivity, comparing relative differences across cells in wages and employment levels between old immigrant and native born workers tells us about the substitutability of the two types of workers. If we see a large relative jump in old immigrant employment but comparatively small changes in their relative wages, for example, we would conclude that old immigrant and native born workers are very substitutable in production - that the old-immigration shock affected immigrant and native born wages similarly. Defining cells as narrowly as possible makes it more plausible that there are no productivity differences between the worker types. Thus, old immigrant and native born agricultural labourers of the same age would reasonably be expected to be working with the same technologies and thus to have the same θ 's. To isolate this cell-level variation, we include separate occupation, age and time effects. This is equivalent to assuming that we can write,

$$\ln\left(\frac{\theta_{N_{kljt}}}{\theta_{O_{kljt}}}\right) = \psi_l + \psi_t + \psi_j + u_{kljt} \quad (11)$$

where the ψ 's are fixed effects and u_{kjt} is an independent error term.¹⁴ Identification then relies on the assumption that relative supplies of native born versus previously arrived immigrants may respond to longer run factors such as the persistent occupation, age and time effects captured in the included fixed effects, but do not respond to the idiosyncratic disturbances to productivity reflected in u_{kjt} .¹⁵

We run equation (11) using real annual earnings and employment numbers from the Census data aggregated to the narrow occupation group level. We drop cells with fewer than 10 observations and weight using the square root of the cell size. The dropping of small or empty cells leaves 1534 occupation by age group cells. In all of our estimations, standard errors are clustered at the occupation x age group level to allow for flexible forms of serial correlation in the error term. When we move to working with new immigrants in the higher nesting levels, we will create instruments to address potential endogenous adjustments by the immigrants, but no such instrument is available for the older immigrants and so we only present OLS results for this regression. In Table 4, we present the key estimated coefficients (the negative inverse elasticities) and the implied elasticities from each stage of the CES nesting. The results correspond to dependent variables constructed using annual earnings. The parallel set of results using weekly wages are very similar and we do not present them for brevity.

Our estimate of the negative inverse elasticity of substitution is -.014 and not statistically significant at any conventional level. This falls in the lower part of the range of estimates from DOP for modern day Germany and OP for modern day US (their estimates lie between -0.01 and -0.06).¹⁶ The corresponding elasticity from our estimates is 72, implying that previously arrived immigrants and native born workers are almost perfectly substitutable in production, as we predicted.

We next move one step up in the nesting hierarchy to consider substitutability between workers already in Canada (both native born and immigrants who arrived at least 10 years before a given Census) and newly arrived immigrants within the same age-occupation cell. To do this, we aggregate (9) to the previous/new immigrant worker level. Thus, for the combination of previously arrived immigrants and the native born the relevant equation is:

$$\begin{aligned} \ln \bar{w}_{ONkljt} = & \ln\left(\frac{\partial F}{\partial L_t}(K_t, L_t)\right) + \frac{1}{\sigma} \ln L_t + \ln \theta_{kt} - \left(\frac{1}{\sigma} - \frac{1}{\sigma_o}\right) \ln L_{kt} + \ln \theta_{klt} - \left(\frac{1}{\sigma_o} - \frac{1}{\rho}\right) \ln L_{klt} \\ & + \ln \theta_{kljt} - \left(\frac{1}{\rho} - \frac{1}{\delta}\right) \ln L_{kljt} + \ln \theta_{ONkljt} - \frac{1}{\delta} \ln ON_{kljt} \quad (12) \end{aligned}$$

where \bar{w}_{ONkljt} is the weighted average of old immigrant and native born wages in the occupation-age cell, with the weights being the proportions of old immigrant and native born workers. A similar equation can be derived for new immigrant workers. Taking the difference between these two equations then yields,

$$\ln\left(\frac{w_{Ikljt}}{\bar{w}_{ONkljt}}\right) = \ln\left(\frac{\theta_{Ikljt}}{\theta_{ONkljt}}\right) - \frac{1}{\delta} \ln\left(\frac{I_{kljt}}{ON_{kljt}}\right) \quad (13)$$

¹⁴Because the narrow occupations are nested in the broad occupations, the inclusion of the complete set of narrow occupation fixed effects would make the inclusion of broad group fixed effects redundant.

¹⁵Ottaviano and Peri(2008) provide an extended discussion of the form for $\frac{\theta_{Nkjt}}{\theta_{Okjt}}$, arguing for modelling it as time invariant since general technological shifts over time are captured in higher levels in the nesting. However, they also note that it is common to allow for the possibility that there is time variation and we follow that tradition by including time effects. It is worth noting that even they must really be assuming some time variation in this ratio because without it there would be no error term in the regression.

¹⁶Both DOP and OP use education groups rather than occupation groups, with 3 education groups and 8 experience groups. Thus, the total number of skill cells is similar to ours when we use broad occupation groups.

Table 3: Negative Inverse Elasticities of Substitution from CES Estimation

	OLS	IV	Implied Elasticity	Elasticity Between	No. of Obs	First Stage F stat
$-\frac{1}{\phi}$	-.014 (.018)		71	Old Immigs & Natv Born	1394	
$-\frac{1}{\delta}$.008 (.018)	.11 (.091)	-	New Immigs & Old Immigs/Natv Born	1258	12.3
$-\frac{1}{\rho}$	-.073** (.012)	-.045* (.019)	22	Age Groups	1127	552
$-\frac{1}{\sigma_0}$	-.10** (.020)	-.10* (.049)	7	Narrow Occ. Groups	423	35.6
$-\frac{1}{\sigma}$	-.28 (.14)	-.32 (.16)	3	Broad Occ. Groups	30	6.7

Notes: The calculated elasticity is based in the IV estimate in all rows except for the elasticity between old immigrants and the native born, where there is no valid instrument.

Specifications for the first two rows include year, age and occupation fixed effects. The remaining rows include year and occupation fixed effects.

Standard errors in parentheses. Standard errors are clustered at the occupation by age group level in the first three rows, and at the occupation levels in the remaining rows.

*, ** correspond to statistical significance at the 5 and 1 % levels, respectively.

As with (10), we will assume that the ratio of productivity terms can be captured by occupation, age and time fixed effects plus an idiosyncratic error term.

While it is common to approach this type of estimation treating immigration as an exogenous supply shock within a given cell, it is also common to be concerned about potential remaining endogeneity. In particular, in our context, we know that immigrants who were explicitly brought in to work in agriculture in the West, were, by the time we observe them in the next Census, in a variety of occupations and locations. Thus, the set of changes in immigrant employment we observe might reasonably be expected to be partly due to endogenous responses to within-cell productivity shocks. One approach taken to this problem by a number of authors, starting with Card(2001), is to instrument for the change in immigrant supplies using a variable based on the notion that immigrants will move to locations within the host country where there are enclaves of people from their country of origin, regardless of the local economic circumstances. We use a variant on this instrument in which we assume that immigrants move to a geographical location based partly on these enclave considerations and then, once there, pursue the most common occupations in that location. That is, we assume that when they search for a job in the location they are attracted to for companionship reasons, their probability of finding a job in a given occupation is proportional to the size of that occupation in that location. Moving to a mining town, immigrants are more likely to find mining jobs than if they moved to Toronto, say.

To construct the instrument, consider an immigrant from source country g , and let the proportion of immigrants from g in location r in year t be given by p_{grt} . Also, let the proportion of all workers in narrow occupation l that is within broad occupation k in location r in period t be given by π_{klrt} . We will predict the number of people from country g in occupation k in year $t+1$ at the

national level as:

$$n10_{klgt+1} = \sum_r \pi_{klrt} * p_{grt} * n10_{gt+1} \quad (14)$$

where $n10_{gt+1}$ is the number of immigrants from country g who arrived in the previous 10 years at the national level. Note that this exercise is done separately by age group with both the π_{klrt} weights and the national level inflows (n_{gt+1}) being age group specific.¹⁷ Finally, we form our instrument for $\ln(\frac{I_{kljt}}{ON_{kljt}})$ as $\ln(\frac{n10_{gt}}{ON_{kljt}})$. The first stage F-stat for this instrument is 12.3, implying that we do not have a weak instrument problem.

The key identifying assumption is that neither the composition of employment by occupation at the start of a period in a location nor the distribution of immigrants from a given country across locations at the start of the period are correlated with the change in productivity within an occupation-age cell at the national level in the ensuing decade. We also require that the general growth in immigration from each country is not correlated with the cell specific productivity shocks. These assumptions would be violated if people from a given source country only immigrated to Canada because there was an increase in demand for a specific occupation in a specific location. So, for example, if all migrants from Lithuania were miners and they all went to Northern Ontario because there was an ongoing increase in demand for miners across decades in that region then the correlation between past concentrations and recent inflows of Lithuanians to the region would partly reflect productivity trends rather than just living situation preferences and the instrument would be problematic. Note that the issue has to do with trends in productivity not levels (i.e., not just the fact that miners do well in Canada) because we include occupation fixed effects. The fact that π_{krt} is calculated for all workers in the location, not just for immigrants from the specific source country, makes it more likely that the requirements for the instrument to be valid are met.

Both the OLS and IV estimates of $-\frac{1}{\delta}$ are positive and insignificant. We interpret this as implying that new immigrants and previously arrived workers are perfect substitutes.

Next, we aggregate the wage equations to the occupation-age level. As with the previous exercise, we can write an expression (which we omit for brevity) for \bar{w}_{kljt} , the weighted average of immigrant and native born wages in the sector-occupation-age cell, with the weights being the proportions of immigrant and native born workers. If we choose an arbitrary age group, $j = 1$, as the base group, we can take the difference between the mean wage of any other group and the base group to arrive at:

$$\ln(\frac{\bar{w}_{kljt}}{\bar{w}_{kl1t}}) = \ln(\frac{\theta_{kljt}}{\theta_{kl1t}}) - \frac{1}{\rho} \ln(\frac{L_{kljt}}{L_{kl1t}}) \quad (15)$$

Since we want to isolate variation within occupation-time cells in order to make it more likely that the relative θ values can be considered to be equal, we run this regression in differences including full sets of separate occupation, age and time dummies. This allows for separate trends for each age and occupation group. We instrument for the employment ratio using the log of the ratio of our $n10_{kljt}$ variable for each age group relative to the base group value for $n10_{kljt}$. The resulting first stage is strong (with an F-stat of 552). The omitted category is workers aged 25 to 39 and we drop that group from the estimation. Our IV estimate of $-\frac{1}{\rho}$ is -.045 and is statistically significant at the

¹⁷The regions are defined as follows. We divide Canada into rural areas, towns, and 13 major cities (Victoria, Vancouver, Edmonton, Calgary, Regina, Winnipeg, Hamilton, Toronto, Ottawa, Montreal, Quebec City, St. John, and Halifax). Towns are all locations with populations over 5000 that are not on the list of major cities, and rural areas are all remaining locations. We define a region as either a major city or a town or rural area within a given province (e.g., towns and rural areas in British Columbia are two of our regional groups). We group Prince Edward Island and Nova Scotia together, resulting in 8 provincial groups. Thus, the total number of regions is $13 + (8*2) = 29$.

5% level. The implied corresponding elasticity of substitution is 22, which is large but considerably below the substitutability among immigrant and native born workers within age x occupation cells.

Next, we aggregate up to the narrow occupation group level to arrive at,

$$\ln \bar{w}_{klt} = \ln\left(\frac{\partial F}{\partial L_t}(K_t, L_t)\right) + \frac{1}{\sigma} \ln L_t + \ln \theta_{kt} - \left(\frac{1}{\sigma} - \frac{1}{\sigma_o}\right) \ln L_{kt} + \ln \theta_{klt} - \frac{1}{\sigma_o} \ln L_{klt} \quad (16)$$

where \bar{w}_{klt} is the average wage in the occupation cell in period t. If we assume that the first three terms on the right hand side of (16) can be captured by occupation and time effects then we can estimate (16) including those effects, using within occupation over time variation. As at the third level, we run the regression in differences and include occupation and time dummies. We instrument for the change in $\ln L_{klt}$ using the log of the $n10_{kljt}$ variable aggregated up to the narrow occupation level. The first stage F statistic is again large at 35.6. The resulting estimate of $-\frac{1}{\sigma_o}$ is -.10 and is statistically significant at the 5% level. The implied elasticity of substitution among narrowly defined occupations within a broad occupation group is 7.

Finally, we aggregate up to the broad occupation group level and estimate the analogous equation to (16) using the same methods. Again, the instrument is based on an aggregated version of $n10_{kljt}$. In this case, we do have a weak instrument problem, with a first stage F-statistic of 6.7. The IV estimate for $-\frac{1}{\sigma}$ is -.32, which is not statistically significant at the 5% level. We also implemented a specification in which we instrumented for the change in the log total employment level using the change in immigration based on Borjas(2003)'s interpretation of immigration as a supply shock. That results in a first stage F stat of 89 and a (statistically significant) estimate for $-\frac{1}{\sigma}$ of -.44. We do not report this estimate in the table because we do not believe in that instrumenting strategy. However, from the combination of those estimates and the OLS and IV estimates reported in the table, we are confident that the estimated inverse elasticity is larger than that for the previous stages. In the simulation exercises in the next section, we use the IV point estimate, -.32, which corresponds to an elasticity of substitution for broad occupation groups of 3.

5.1.2 Counterfactual Exercises

The elasticities estimated to this point are of limited intuitive value on their own but can be used with the equation for native born wages, (9), to investigate the impact of immigration on native born wages. As OP point out, inspection of (9) reveals several channels through which an immigrant shock in a specific occupation by age cell can affect the wage of any native born worker defined by his broad occupation group, his narrow occupation within that group, and his age. First, the fact that labour of different types is imperfectly substitutable implies that all workers benefit from an increase in the total number of workers in the economy. This is captured in the second term on the right hand side of (9), which is positive. The fourth term on the right hand side captures the effect of increased immigration within the same broad immigration group. This effect will be negative if the elasticity of substitution among broad groups is less than that among the narrow groups within broad groups (as seems likely). This, and the overall size effect captured in the second right hand side term, are the channels through which immigration in the same broad occupation group but a different narrow occupation group can affect a worker's wage. Similarly, the sixth term allows us to capture the effects of immigration in the same narrow occupation but different age group and the eighth term reflects the direct effect of immigrants entering in the same narrow occupation and age group. The last two quantity terms in the expression relate to previous arrivals and the native born and so won't change with immigration.(Ottaviano and Peri(2010)).

Our approach to quantifying the impact of immigration on native born wages is to construct a set of counterfactual wage distributions. In particular, we focus on the 1921-31 decade and

construct counterfactuals corresponding to: 1) no immigration between 1921 and 1931 (in order to match the most common counterfactual considered in the previous literature); 2) no general or agricultural labourer immigration in that decade (in order to examine the inequality implications of the dominant, low skilled component of the inflow); and 3) no immigration of professionals in that decade (as a point of comparison for 2)). To construct the first counterfactual, consider generating a fitted wage for 1931 using equation (9) but with the number of recent arrival immigrants subtracted from each relevant labour measure on the right hand side. That is, form a new version of L_{kt} as $L_{kt}-I_{kt}$ (where I_{kt} is the number of recent immigrants in broad occupation, k), a new version of L_{klt} as $L_{klt}-I_{klt}$, and so on. Importantly, this counterfactual is based on the number of new immigrants observed in the 1931 Census and so already nets out any effect that emigration would have in reducing the impact of initial immigration inflows. Since the productivity terms (the θ 's and $\ln \frac{F}{L_t}(K_t, L_t)$) are assumed not to change with immigration, we can construct the counterfactual wage in 1931 as,

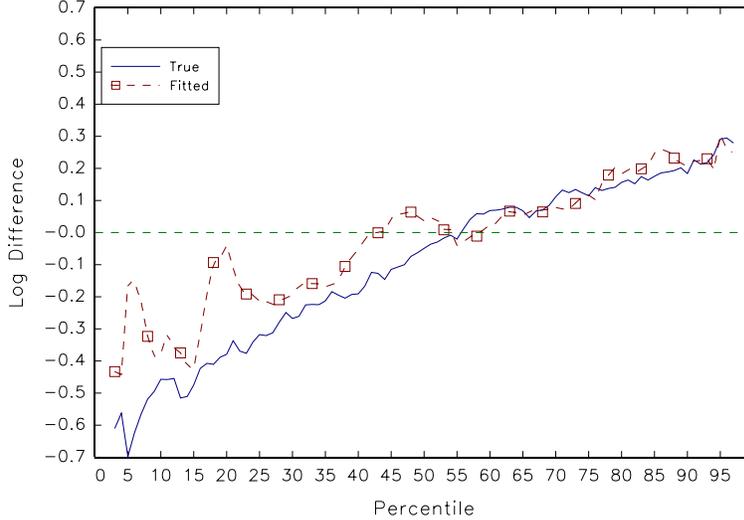
$$\ln w_{Nkljt}^c = \ln w_{Nkljt} - \frac{1}{\sigma} \ln I_t + \left(\frac{1}{\sigma} - \frac{1}{\sigma_o}\right) \ln I_{kt} + \left(\frac{1}{\sigma_o} - \frac{1}{\rho}\right) \ln I_{klt} + \left(\frac{1}{\rho} - \frac{1}{\delta}\right) \ln I_{kljt} \quad (17)$$

We then construct a complete counterfactual distribution by creating a set of "phantom" people for each occupation x age cell. That is, if a given cell has n_{klj} native born workers in it in 1931 then we generate n_{klj} phantom people all with the fitted wage, $\ln w_{Nklj}^c$. Doing this for each cell, we obtain a dataset that has the same size and distribution of workers across cells as the true 1931 data but with fitted wages that reflect the estimated elasticities and the counterfactual quantities.

To capture the impact of not having immigration, we need to compare this counterfactual distribution to the actual distribution. This requires forming a version of the true distributions that eliminates within occupation by age cell wage variation, just as the counterfactual exercise does. For 1931, we do this by again constructing a set of phantom workers for each cell but in this case we assign them the actual average wage for the relevant cell in 1931. We do this for both 1921 and 1931 then calculate the percentiles of these generated datasets and take the log difference. In Figure 7, we plot the resulting difference line (which we label the "fitted" difference) along with the differences in log percentile between these years from the actual dataset (replotted from Figure 1 and labelled here as "True"). The result shows that removing the within-cell variation does not substantially alter the picture of increased earnings inequality in the 1920s. In terms of the decomposition discussion in the previous section, the fitted line shows the effects of changes in the wage structure that arise only because of changes in between group wage differentials. The fact that it lies above the actual line below the median implies that both between group differential changes and changes in dispersion within groups explain the increase in inequality in the lower part of the distribution. On the other hand, above the median, the fitted and actual lines are coincident, implying that the earnings changes in that part of the distribution arise only because of within group dispersion.

In Figure 8a, we re-plot the fitted line from the previous figure along with the difference between the log percentiles of the 1931 counterfactual distribution with no recent arrival immigrants and the 1921 fitted distribution. The result is a small implied immigration impact across the distribution. This can be seen more clearly in Figure 9, where we plot the difference in log percentiles between the counterfactual distribution and the fitted 1931 distribution. This difference provides our key estimate of the net impact of immigration on the earnings structure in this period. Because the number of immigrants we remove from the 1931 labour force to construct the counterfactual come from the Census, they reflect the adjustments made by immigrants within Canada. Given our earlier arguments that the occupational distribution of new immigrants is similar to that of the workers already in Canada, taking them away from the distribution implies small changes in the distribution.

Figure 7: Actual and Fitted Differences in Log Percentiles, Male Earnings, 1921–31

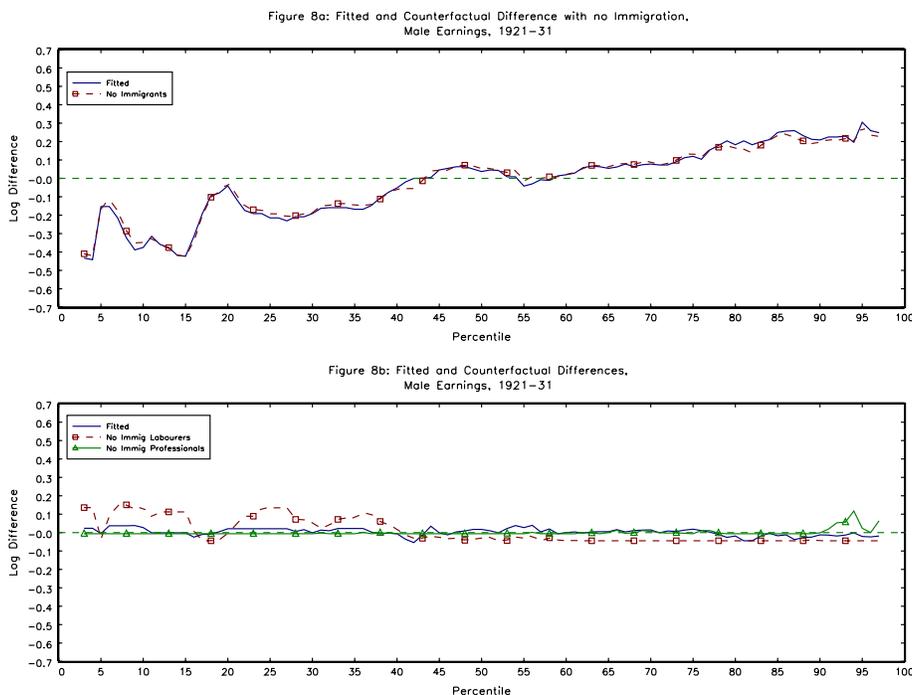


But the occupational distributions of new arrivals versus settled workers are not identical and the small impact partly reflects general equilibrium adjustments. We return to a measure of the relative importance of those various channels for one group of workers below. In Figure 8b, we plot differences corresponding to the other two counterfactuals: eliminating only recent immigrant labourers and only recent immigrant professionals. The implied effects of the labourers alone are substantial. Below the 40th percentile, removing the labourer inflow would have resulted in an average increase in native born real earnings of nearly 10 percentage points. Again, this can be seen clearly in the plots relative to the true 1931 distribution in Figure 9. There, one can also see the spillover impact of removing immigrant labourers on wages in other occupations. In particular, real earnings in occupations with above median wages would have fallen by approximately 5%. This highlights the arguments in OP that one needs to look beyond narrow occupation groups to measure the full impact of immigration. It also implies that the impact of low skilled immigration alone was to substantially increase inequality. The growth in the difference between the log 90th and 10th percentiles in the 1920s would have been reduced by nearly a third if no immigrant labourers had been admitted. The growth in the 50-10 differential would have been reduced by over 50%. The other counterfactual shows the impact of not admitting recent professional occupation immigrants. Removing that inflow would have had substantial positive effects on wages above the 90th percentile but only small negative effects elsewhere in the distribution.

We can use different predicted impacts on farm labourers under various scenarios as a means of highlighting the extent of the adjustment channels we have discussed. In the Appendix, we present estimates of the main specification from Borjas(2003) in which we regress decadal changes in native born wages on the proportional increase in farm labour that is due to recent immigration.

¹⁸ This specification can be given an interpretation as showing estimates of the elasticity of demand

¹⁸That is, we estimate the regression, $\Delta \ln(w_{Njkt}) = \alpha_0 + \alpha_1 \frac{\Delta E_{Ijkt}}{E_{Njkt-1} + E_{Ijkt-1}} + \Delta u_{jkt}$, where $\Delta \ln(w_{Njkt})$ is the change in native born log earnings in occupation j and age group k , ΔE_{Ijkt} corresponds to the number of recent immigrants who were in that occupation and age category in the time t (e.g., 1931) Census, and $E_{Njkt-1} + E_{Ijkt-1}$ is the total number of workers (native born and immigrant) who were working in that occupation by age cell in the $t-1$ Census (1921 in this example). Note that the regression is based on Census data and thus the immigrant employment numbers are after immigrant adjustments.

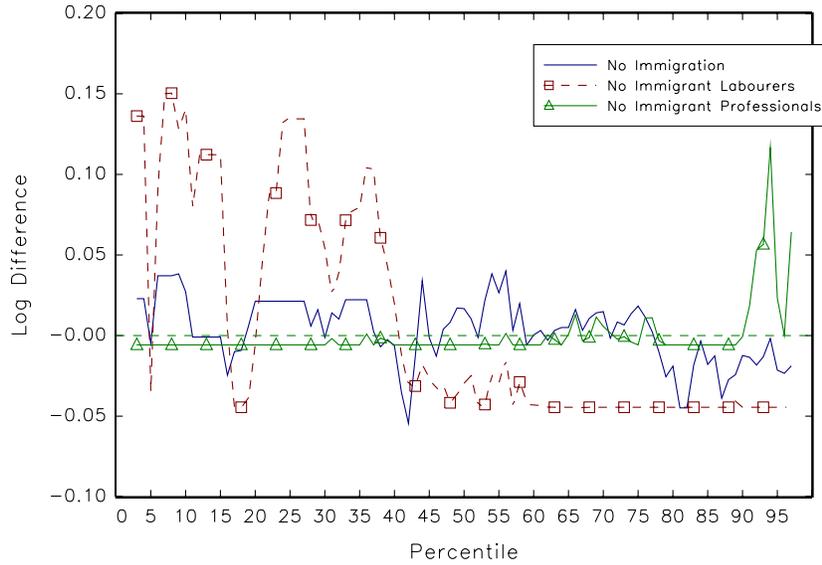


if immigrants and native born workers were perfect substitutes and each occupation represents a separate labour market. That is, the estimated effects can be seen as reflecting immigrant effects on earnings when there are no further general equilibrium effects. Using the narrow occupation groupings in order to match with our more structural estimation, the estimated coefficient on the proportion of workers who arrived in the previous decade is -0.35 . If all the immigrants who arrived in the 1920s had remained in Canada and in their intended occupation, the inflow of new immigrants over the 1920s would have implied a 158% increase in the farm workforce and thus a 42% decline in worker earnings in that sector.¹⁹ However, the actual number of farm workers in 1931 who were immigrants who had arrived in the previous decade (i.e., the number after accounting for adjustments through re-migration and occupational switching) amounted to an 39% increase in the farm workforce, and thus would imply an 13% decline in farm worker earnings in the absence of other general equilibrium effects. Finally, making use of our production function parameter estimates and the full set of immigration inflows, our first counterfactual implies that the total immigration inflow implied a 4% drop in farm worker earnings. This last reduction (from the 13% effect to the 4% effect) represents the impact from taking account of general equilibrium adjustments in the economy.

The counterfactual with no immigration implies that immigration had only limited impacts on the location of the native born earnings distribution. This fits with findings in Pope and Withers(1994) for Australia and Greasley and Oxley(2004) for Australia. Taylor and Williamson(1997) and Hatton and Williamson(1998) generate much larger wage impacts from immigration using CGE models with capital held fixed, but those impacts are reduced by as much as 75% once capital is allowed to follow labour. Green and Sparks(1999) find evidence of such endogenous capital inflows

¹⁹For this calculation, we adjust the total inflow for the fact that the inflow statistics include farmers as well as farm labourers. To do this, we multiply the inflow in the "Farmers and Farmer Labourers" category by the proportion of the sum of farmers and farm labourers who were farm labourers among immigrants who had arrived in the previous decade in the 1931 Census (0.65). For this calculation and the ones that follow, we use the number of males employed as farmer labourers in the 1921 Census as the base.

Figure 9: Differences in Log Percentiles Between Counterfactual and Fitted Distributions, Male Earnings, 1931



for Canada before WWI using a cointegration methodology. A finding of immigration having small impacts on native born wages also fits with the literature examining recent immigration (e.g., Ottaviano and Peri(2012)). Methodologically, Dustmann and Preston(2011) argue that the type of method we use understates immigration impacts to the extent that immigrants have a certain set of skills on paper but actually end up competing for jobs with less skilled native born workers. In our case, though, where we match immigrants to the native born in terms of the occupations they actually work in and where immigrant and native born occupational distributions are very similar, this problem is unlikely to arise.

Impacts on inequality, however, may differ from impacts on the average wage level. Abad(2013) finds substantial impacts of immigration on inequality in Latin America in the nineteenth century, which does not fit with our counterfactual result that total immigration had small impacts on inequality. However, Abad(2013) measures inequality using the ratio of land rents to the unskilled wage, and we do find in our second counterfactual that the substantial unskilled labour inflow in this period reduced earnings in the lower part of the distribution. We conclude that large inflows in specific occupations can have substantial impacts on the wages in that occupation and in closely substitutable occupations but that can be mitigated by positive spill-over effects due the arrival of a wider set of immigrants. In our case, the effects of those other immigrants implies near zero net effects of immigration even in low skilled occupations.