

Does Completion Matter: Measuring the Disparity in Wage Outcomes between Completers and Non-Completers of Apprenticeship Programs in Arkansas

Aimee Jahnke, Department of Workforce Development, Wisconsin
Corey Goodrich, Department of Workforce Development, Wisconsin
Dan Sharp, Madison Area Technical College, Wisconsin
Dominic Robinson, Department of Workforce Development, Wisconsin
Heather Saco, Division of Information Systems, Arkansas
Lesia Edwards, Department of Commerce, Arkansas
Priyanka Bejugam, Division of Information Systems, Arkansas

1 Introduction

Apprenticeship is a training strategy that prepares people for skilled employment by conducting training in bona fide and documented employment settings. The content of training, both on-the-job and related instruction, is defined by the industry.

Apprenticeship is post-secondary education like a college or university. But there's a big difference. Registered Apprenticeship (RA) combines paid on-the-job training under the direction of a mentor with classroom instruction and results in an industry-recognized credential. There is also a built-in wage progression in the contract, so the apprentice receives most of their training on the job while working for an employer who pays a good wage.

What makes it registered is that there are government standards and registration processes for it to be considered a RA program with defined standards. In apprenticeship, 90% of apprenticeship training is learned on the job. There are written standards that govern on-the-job training – they spell out to a T – what the apprentice must learn throughout the apprenticeship. The work must be supervised by a skilled worker who trains the apprentice – passing on the skills and knowledge of the trade.

A critically important feature is that to start in RA, you must be an employee first, then enroll in RA program through a sponsor; this may be an employer, employer organization, labor-management organization, or committee. The sponsor oversees and evaluates the work and training of the apprentice to ensure that the apprentice will become a skilled journey worker. The sponsor is usually the apprentice's employer in all but the construction trades.

In construction trades, the employer is not the sponsor. Because apprentices in the construction trades tend to move from one employer to another, depending upon available work, the construction trades apprentice sponsor is a formal local committee overseeing the apprentice's training. The local committee assigns the apprentice to an employer. They are the ones that recruit and set minimum standards to get in, and once you are enrolled, there is a process that results in a contract outlining what to expect from both the Sponsor and the Apprentice. All three parties sign that, the sponsor, the apprentices, and the government agency.

The Sponsor is most often responsible for most of the costs, including the pay and expenses for experienced employees to help train and mentor the apprentice in the OJT. While the government agency is responsible for standards in place, compliance with the standards, and contracts. They also provide operational tools to track programs, metrics, and contract details. The government agency can also fund incentives to support the use and access to this training method.

What motivated your research project?

The project's objective was to enhance the team's understanding of the various tools and techniques used in Work Based Learning programs. We also wanted to develop a framework that would help improve the effectiveness of the Registered Apprenticeship Program. As a result, the main goal of this project was to build capacity and develop a framework for bringing together the various elements of the Registered Apprenticeship Program into a single language that can be used across different program types and state lines. We also explored the possibility of developing new tests and measures that could help improve the effectiveness of RA.

Many current measures are suitable for the job they're designed to do, but they do not provide enough flexibility to ensure that contracts are enforced. The lack of consistency across state programs and the tools that end with the agreement also did not help improve the effectiveness of the program. One of the most important factors employers and program administrators need to consider when implementing information versus measures is the audience they are trying to reach. For instance, this may be an appropriate measure if the audience is focused on Government Investment.

What is your research question?

Does Registered Apprenticeship completion impact wages?

Suggestive evidence shows that completion positively and significantly impacts wage outcomes. Statistically and Economically Significant: 40% higher earnings per year on average. The other factors tested were gender, race/ethnicity, age, location of work, and reported education at entry. We found that being white and non-Hispanic did affect wages. White non-Hispanic individuals in our sample, holding all else constant, made 23% higher wages than their counterparts. There were more profound questions raised, as our descriptive analysis suggests completers start at higher levels of salaries and exhibit similar wage growth as non-completers.

We saw that these results were similar across apprenticeship occupations. Among types that have a large enough sample size were the Plumbers and Electricians; we see identical wage differences between completers and non-completers.

If we did not find a match in the UI wage file, we removed them for this analysis. We cannot infer their wages were 0, as we do not know if they worked out of state or in uncovered employment like self-employment. There was a significant difference in the percent of people we had to remove between completers and non-completers. We had to remove more non-completers. So, although we cannot infer that these people do not have a wage if we did not find a match, the size of the difference can suggest that non-completers are more likely to be not working; this would be an area that needs additional study.

We cannot infer that they do not have a wage if we do not find a match. But the size of the difference can suggest that they are more likely not to be working; this would be an additional study area. So, if you are white, non-Hispanic, and holding all else constant, you earn 23 % more.

2 Data

For our analysis, we utilized apprenticeship data from the RAPIDS universe which allows us to identify key characteristics of apprentices within the state of Arkansas. Of particular interest for this analysis, the RAPIDS dataset lets us identify the status, type, and length of each apprentice's program as well their age, race, ethnicity, gender, and education at program entry. The RAPIDS data was then merged to quarterly unemployment insurance (UI) wage data utilizing a crosswalk of RAPID apprentice numbers to Social Security Numbers (SSN). It is important to note that by utilizing this crosswalk the majority of the existing apprentices were dropped from our sample due to no identifiable SSN¹.

With the UI wage data, we can observe wage data for apprentices within the state of Arkansas as far back as 2008; however, UI wage data has significant data limitations. Chief among these limitations, is that missing data cannot be inferred to be a null since UI wage data can be missing if the individual works out of state, either remotely or in-person (e.g., the individual moved out of state), or if the individual works in a UI uncovered position, such as an independent contractor. This data missingness is significant since the majority of the sample that remains after the crosswalk has incomplete or no wage data. Further, since the data just shows the quarterly sum of wages, we are unable to make any judgement on the number of hours worked relative to wages.

2.1 Cohort Design

The cohort that is utilized in this report is designed to allow for a comparison of wage outcomes between completers and non-completers². As such, any apprentice that doesn't have a program status of completed or canceled, such as a currently enrolled apprentice, is dropped from our cohort. Due to programs varying in length, we utilize the data elements that collect information on the start date and the expected program end date to calculate the expected program length (EPL). This then allows us to align the varying program lengths to a single measure such that we can perform analysis on the event of completion, as well as a similar post-completion measure. Our EPL measure is defined as the following:

$$100\% EPL_i = \text{Expected Program End Date}_i - \text{Program Start Date}_i$$

Where 100% refers to the date that the program was expected to be completed. In addition to EPL, we the calculate 150% EPL to allow for a post-completion analysis. Where the 150% is a set time after expected completion date that will vary based on the program length. For instance, a 150% for a two-year program would be at the mark one year after program exit, while for a four-year program it would be two years after program exit.

$$150\% EPL_i = 1.5 * (\text{Expected Program End Date}_i - \text{Program Start Date}_i)$$

To avoid complexities in our data analysis, particularly with the impacts of COVID-19, we restrict our cohort to apprentices that reached 150% EPL³ at any time between Q1 2016 and Q4 2019. This timeframe

¹ While our analysis doesn't analyze if there were any inherent differences between those that are matched and those that aren't, any further analysis should verify that this crosswalk procedure isn't biasing the result towards a select group of apprentices.

² As will be expanded upon in the caveats section of this report, this comparison is flawed. Due to time limitations of the authors, as well as data limitations, the proposed analytical framework and cohort design should be seen as a very early first step towards creating a generalizable method of measure wage outcomes of apprentices.

³ 150% EPL is used rather than 100% EPL so that our post program completion analysis is not effected by COVID-19 impacts in 2020 and 2021.

is beneficial because the years in our period of study are relatively economically homogenous and there is a large enough sample of apprentices for detailed wage analysis (n=2,950).

2.2 Analytical Model

The proposed analytical framework of this report uses an event study approach to analyze the effect that completion, our main variable of interest, has on the wage outcomes of apprentices. Our report is broken up into two main types of data analysis around the idea of completion, the first being a descriptive analysis of the individuals within our cohort and the other being a multivariate regression to help tease out the effect that completion has on wage outcomes. The multivariate regression is designed in the following way, with analysis being done at 100% and 150% EPL:

$$\begin{aligned}\ln(100\% \text{ EPL wage}_i) &= \alpha + \gamma \text{Completion}_i + \beta X_i + \theta_i + \delta_i + \varepsilon_i \\ \ln(150\% \text{ EPL wage}_i) &= \alpha + \gamma \text{Completion}_i + \beta X_i + \theta_i + \delta_i + \varepsilon_i\end{aligned}$$

In this regression model our dependent variable, wage, is taken as a natural log. We utilize a natural log due to wages having skewed distribution and because it simplifies the interpretation of our results since a marginal change in our independent variables can be interpreted as multiplicative change in wage. On the right-hand side of our equation, θ_i is defined as time fixed effects and δ_i is occupation of enrollment program fixed effects. Among our vector of controls, X_i , we control for the gender, age, race, ethnicity, metro status⁴, and education of the apprentices. It is important to mention that the education variable utilized in this analysis is education at program entry. This is a limitation of the available data, and as a result if any post apprenticeship enrollment educational attainment occurred, we are unable to control for these changes.

3 Key Findings

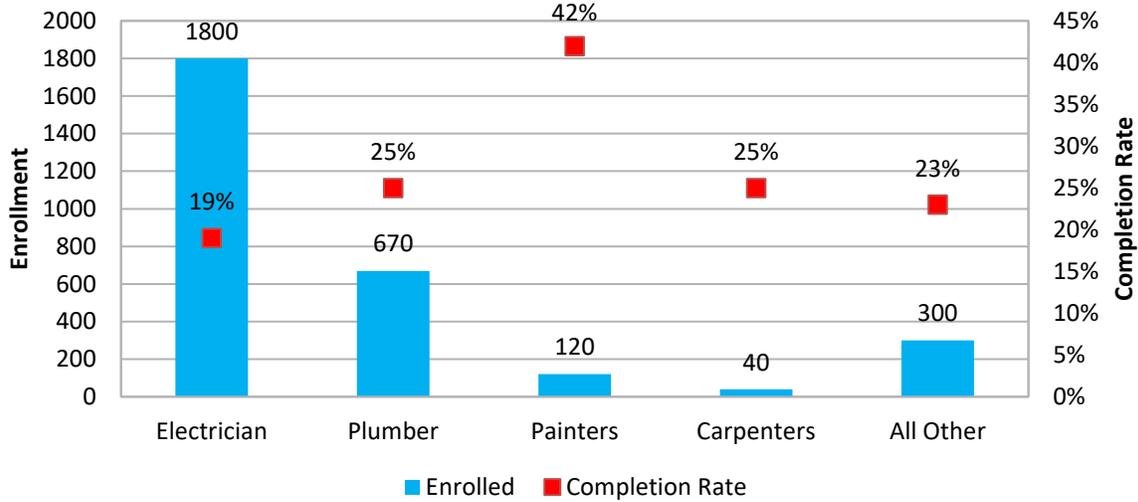
The report will proceed as follows, we will first discuss the descriptive analysis of our cohort to provide context on the individuals within our sample, as well as to provide possible areas of further research. We will then proceed into our regression analysis, which helps contextualize the descriptive analysis while building additional confidence in our results since we are able to ascertain the suggestive evidence on the effects of the covariates.

3.1 Descriptive Analysis

To help provide context to the individuals within our sample, we evaluated what program individuals within our sample enrolled in, and also what the rate of completion was for these programs. Figure 1, which displays our findings, shows a high concentration of enrollees within a few traditional apprenticeship programs, electrician, and plumber. What is most striking from this figure is that apprenticeship programs suffer from very high attrition rates. Regardless of the program, we find that the majority of individuals do not end up completing the program with this being particularly pronounced in the electrician apprenticeships. While we do not delve into what is causing this high cancellation rate, this should be explored further in additional research since the findings would be valuable for policy makers, the figure helps to illustrate that only a minority (n=650) of individuals in our entire cohort (n=2,950) actually complete their program.

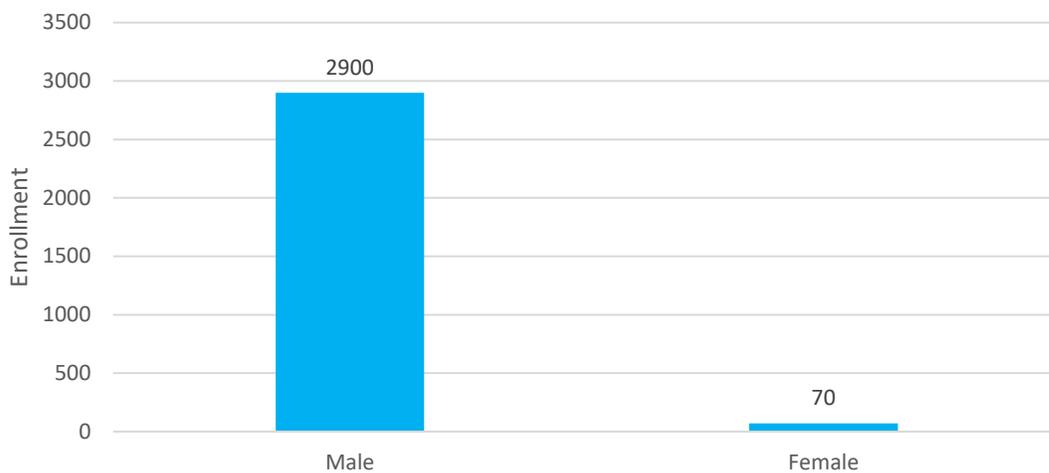
⁴ Metro status is based off the location of the employer with the highest wage for that individual.

Figure 1: Enrollment and Completion Rate by Apprenticeship



Continuing our exploration into variations in enrollment and completion rates within our cohort, we next analyze enrollment and completion rate by gender, as partially shown in figure 2. As can be seen, our cohort is largely made up of male apprentices, with just a small section, roughly 2% of our cohort, being female. By having such a small section of our sample be female, we ultimately had to cut the completion rate from the figure. The reason behind this, is that due to export requirements on the data the rounding of sample sizes meant that the comparison rate of females was drastically different from the actual value. In particular, it would have shown a significant difference in the completion rate for females compared to males, however, this is not the case as they exhibit similar rates of completion. Nevertheless, what figure 2 highlights for policy makers is the need for targeted outreach to females. The enrollment rate of females within apprenticeships is very low even at the national level (12.5%), however, it is particularly pronounced in our cohort.

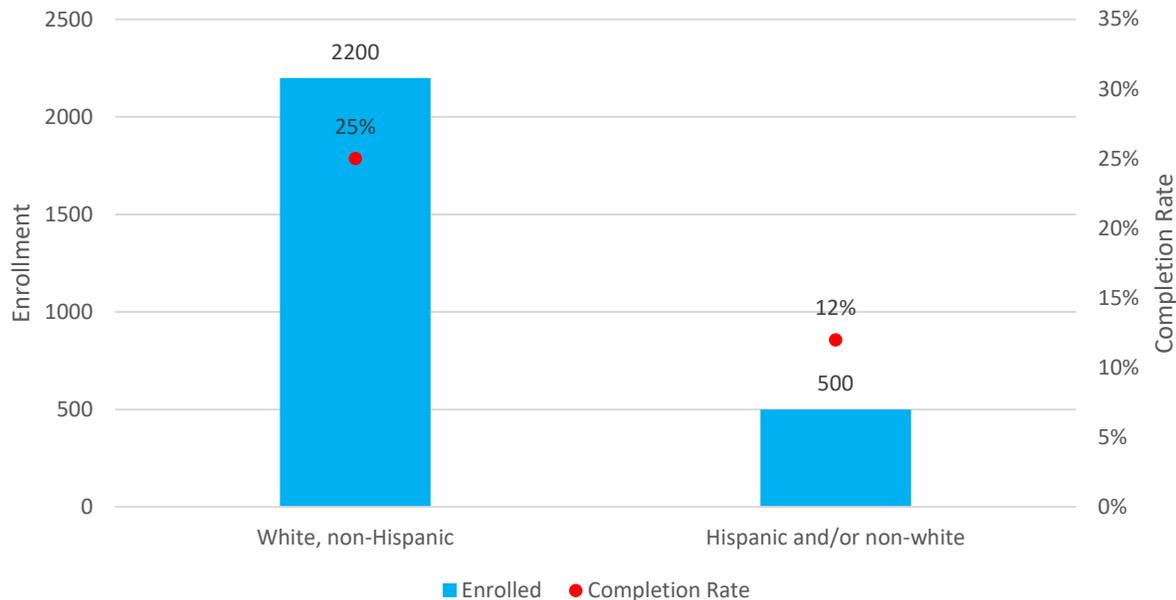
Figure 2: Enrollment by Gender



Shifting into enrollment and completion rates by race and ethnicity, we find additional under representation in our cohort, although not at the scale of gender. Particularly, Hispanic and or non-white individuals make up 19% of our cohort, however, we would expect that they would comprise roughly a third of our cohort given their share of their population in Arkansas. What is also striking from this figure is that Hispanic and or non-white individuals complete programs at half the rate of white non-Hispanic individuals. This should immediately raise questions as to why we see lower completion rates among Hispanic and or non-white

individuals. Again, the aspect of why is not explored in this report since it is beyond the scope, however, this is a point that should be explored more fully. Understanding the why here, will help to inform policy makers on how to increase completion rates for minorities and also how to increase the number of minorities enrolled in apprenticeships.

Figure 3: Enrollment and Completion Rate by Race and Ethnicity



Expanding our descriptive analysis into the wages of our cohort, we utilize data on the expected completion date of apprentices to align our cohort to a time (T0), which is the quarter of expected completion. With a comparable point of time, we then calculate the average wage of completers and non-completers relative to T0, where T-1 refers to one quarter before the expected completion quarter and T+1 refers to one quarter after the expected completion quarter. From here we then take a rolling average of wages to smooth our results from the two quarters before, the reference quarter, and the two quarters after. Due to data missingness in wages, and the infeasibility of assigning nulls to missingness, the average value for each quarter is based off of the average wages of the sample population that have wages for that quarter. This means that individuals can drop in and out of our averages depending on the availability of their wage data. We do not expect this to overly impact the results given that we are working with a large sample of individuals in both groups so the impacts should be minimal. Then finally, the final step that is taken in our descriptive analysis of wage, was to convert our quarterly wage data into yearly values since yearly values are easier to interpret.

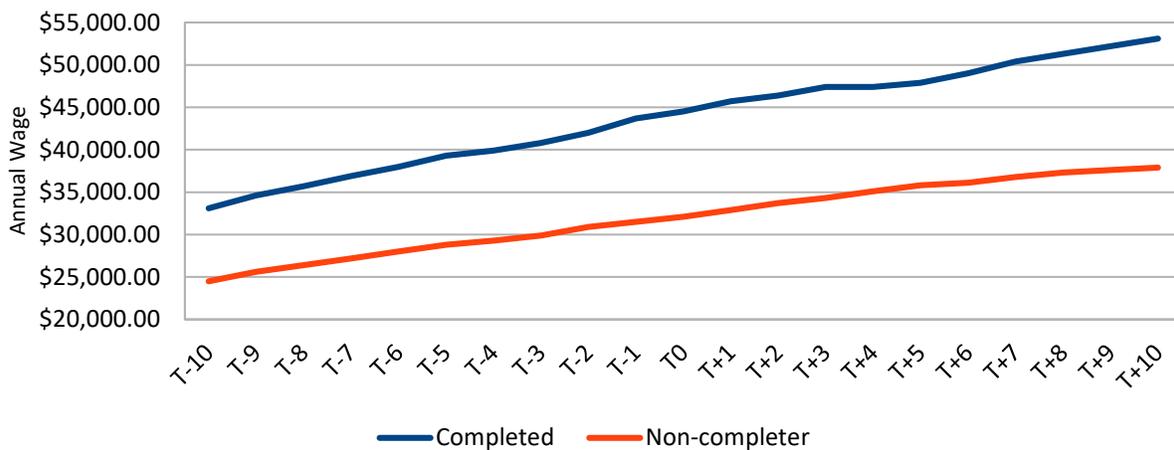
Utilizing our created dataset, figure 4 shows the average wages of completers versus non-completers in our cohort. What is immediately striking from this figure is not that there is a difference, this was expected, but that this difference existed from ten quarters before completion. We also don't really see an impact at the time of expected completion, as one might expect⁵. Our results suggest that while the individual act of completing the program doesn't increase the individuals wage, being an apprentice that completes is associated with higher wage outcomes. Furthermore, more detailed analysis of our results shows that there is a divergence of wage outcomes as time progress, but interestingly, the rate of wage growth is similar in both groups.

⁵ Given that we smoothed the data, the act of completion might be dampened by the averaging, however, prior to data smoothing we saw little to no impact at the time of expected completion.

There are several important caveats to this figure, first while we cannot infer that a missing value is a null, we do find that non-completers were over two times more likely to have missing values. There are couple possible explanations to this, such as more non-completers work out of state or non-completers are more likely to work in non-covered occupations, and while we cannot rule out these possibilities, the significant difference in wage missingness suggests that non-completers are more likely to be out of the labor force.

The second caveat to this figure is that going back T-10 quarters will predate enrollment into programs for some of the apprentices. While we do not believe this will significantly impact our results, since the majority of our individuals are in programs that are four plus years long, further analysis should check the robustness of the results shown here. The last caveat to be mentioned is that for completers the expected completion date is not always the same as the actual completed data. Expected completion date was chosen so that we could have direct comparison to non-completers, however, as will be expanded upon further in the caveat section this comparison is flawed. Future analysis on completers should instead set T0 of completers to the actual date of completion to test if our results are robust.

Figure 4: Comparing average wages of completers and non-completers relative to expected completion date



As a final step in our descriptive analysis, we created a table showing a comparison of means (i.e., balance table) between completers and non-completers in our cohort. What table 1 shows is that completers and non-completers are very distinct subpopulations of our cohort. On nearly every measure a statistically significant difference between the completers and non-completers exists. This table really brings to the forefront, the issues in doing a direct comparison between completers and non-completers. Namely, completers and non-completers are not different based just on the act of completion, but a whole host of other covariates. Each of these covariates, as well as ones that are missing due to no data (which will be discussed further in the caveats sections), can play a pivotal role in wage as well as completion.

Table 1: Comparison of means between completers and non-completers

Covariates	Completers (1)	Non-completers (2)	Difference (1)-(2)
<i>Panel A: Annual Wage</i>			
Wage at 100%	\$44,400.00	\$32,600.00	\$11,800***
Wage at 150%	\$50,000.00	\$36,800.00	\$13,200***
<i>Panel B: Time variant demographics</i>			
Age at 100%	31.86	33.41	-1.55***
Age at 150%	33.50	35.13	-1.63***
Metro at 100%	72.00%	67.00%	5%**
Metro at 150%	73.00%	66.00%	7%***
<i>Panel C: Time invariant demographics</i>			
Male	98.00%	98.00%	0%
White non-Hispanic	89.00%	79.00%	10%***
Less than high school degree	5.00%	11.00%	-6%***
High school	81.00%	78.00%	3%**
Some college	13.00%	11.00%	2%
Observations	650	2300	

3.2 Regression Analysis

As a final part of our analysis, we conduct a regression analysis to tease out the effect that completion has on wage outcomes. Table 2 displays our multivariate regression results. We run four different regressions, two of which are at 100% EPL and the other two at 150% EPL. These two different regressions allow us to measure the impact completion has on wages at the expected date of completion, and the effect it has beyond the event of completion. We find suggestive evidence that completion has a positive and significant impact on wage outcomes of apprenticeship enrollment and that this impact is also economically significant. We see similar wage impacts in both our naïve regression as well as the different EPL variables, suggesting that our results are robust to additional cofounders. The estimated impact of completion on wage is around a 40% increase in wage outcomes, when all else is held constant.

Beyond completion rate, we find suggestive evidence that being white non-Hispanic is associated with higher wage outcomes, however, the impact is dampened over time. At 100% EPL, the estimated impact of being white non-Hispanic is a 23% increase in wage, however, at 150% EPL the estimated impact shrinks to a 15% increase in wage⁶.

What table 2 helps to show is that when considering other factors, we still find that going through a program and completing it is a significant determinant on future wage outcomes. This is important from a policy perspective because it shows that apprentices do play a significant role in the financial health of the person and the community. Couple this with the previous descriptive analysis on low completion rates and this result helps to show the need to bolster the number of individuals that finish their program.

⁶ It is important to note that white non-Hispanic at 150% just misses out on the traditional levels of significance.

Table 2: Multivariate regression

	Wage at 100% expected program length		Wage at 150% expected program length	
Completed	0.44*** (0.04)	0.41*** (0.04)	0.38*** (0.04)	0.38*** (0.04)
Male		0.22 (0.26)		0.26 (0.26)
White non-Hispanic		0.23** (0.09)		0.15 (0.09)
Age		-0.01 (0.02)		-0.02 (0.02)
Age^2		0.00 (0.00)		0.00 (0.00)
Metro (place of work)		-0.07 (0.05)		-0.07 (0.05)
Less than high school degree		-0.13 (0.11)		-0.12 (0.11)
Some college		0.10 (0.07)		0.08 (0.07)
Observations	1500	1300	1400	1300
Time fixed effects	Yes	Yes	Yes	Yes
Occupation fixed effects	Yes	Yes	Yes	Yes
R-squared	11%	14%	11%	13%

4 Caveats

As stated earlier in the report, this analysis should be seen as a first step towards a generalizable way of measuring program outcomes. One of the main caveats to our results is that our comparison group was ultimately flawed. This is because there are likely inherent differences between those that complete and those that don't, that we are unable to control for. Since we are unable to control for important cofounders between both groups our regression analysis will be affected by omitted variable bias. Future analysis should utilize machine learning to develop a more valid comparison group to completers.

Beyond difficulties in comparing completers and non-completers, our analysis was significantly limited by not being able to associate missingness in our wage data to nulls. As discussed earlier, this likely biases our results downwards, since we anticipate that non-completers are much more likely to have actual nulls than the completers. Also mentioned earlier is that we do not have education data post enrollment, so we are unable to measure any changes in education post enrollment in a program which is of particular importance for the non-completers who might be more likely to enroll into a 2-year or a 4-year degree program after exit.

The last caveat to mention is that we do not have occupational information of individuals in our cohort. This means that we can only make assumptions that those that complete their apprenticeship stay within that occupation. While for non-completers, we cannot make any judgement based on their occupation. This is an important area of future research because the researchers can delve into what type of occupations the non-completers end up going into when they cancel their apprenticeship.

5 Possible Extensions

There are several points of extension from this report. First, as mentioned in the previous section a more valid comparison group to completers should be created. Some of the existing literature in the field has utilized machine learning to create a more balanced comparison group to completers, which as table 1 shows non-completers are not. By being able to create a more valid comparison group, there are less concerns over omitted factors, which builds confidence in the results shown.

Secondly, future research should work on redesigning and refining the regression model proposed in this report. The purpose of the regression model in this report was not to design a model that could be used in an academic publication, but instead to show how regression analysis can be utilized in evaluating apprenticeship programs. As a result, future research should design a model that utilizes more complete data and better controls for additional cofounders that can pass a peer-review process in an academic publication.

Lastly, future research should explore the following questions. What stage do non-completers leave, and why do they leave? Do the characteristics of the employer, training provider, or committee characteristics influence completion rates? Do we find that occupations of non-completers drastically change after their apprenticeship is canceled or do they stay within the same or similar field?

6 References

https://www.apprenticeship.gov/sites/default/files/women-in-apprenticeship-fact-sheet_0.pdf
<https://www.census.gov/quickfacts/AR>