Problem 1

Calculate the following test-statistic:

$$T = \frac{\bar{X} - \bar{Y}}{\sqrt{Var(\bar{X} - \bar{Y})}} \tag{1}$$

where the pair (X_i, Y_i) correspond to a pair of husband and wife's respective earnings. We can calculate $Var(\bar{X} - \bar{Y}) = \frac{Var(X_i - Y_i)}{n} = \frac{Var(X_i) + Var(Y_i) - 2 \cdot Cov(X_i, Y_i)}{n} = \frac{Var(X_i) + Var(Y_i) - 2 \cdot \rho \cdot \sqrt{Var(X_i)} \cdot \sqrt{Var(Y_i)}}{n}$. We are given the sample versions of all of these quantities in the question (i.e. we use sample variance as an estimate of the population variances). Our null is $H_0 : \mu_X = \mu_Y$. First, do a one-sided test. I will test that male earnings are greater than female earnings: $H_1 : \mu_X - \mu_Y > 0$. Plugging in, we get $T = \frac{55{,}651 - 45691}{9.8629} = 1009$. Because n is large, this statistic is approximately normally distributed. For a one sided test with $\alpha = 0.05$, the critical value is 1.64. Therefore we reject the null. For the two sided test, we have $H_1 : \mu_X \neq \mu_Y$. The test statistic is $|T| = 1009$, and critical value is 1.96. Therefore, we reject again.

Problem 2

a) (i) Here we have independent samples. I will perform a test of the equality of the means between immigrants and native Canadians assuming equal population variances (seems like a reasonable assumption here; it is okay if you assume different population variances and calculate the corresponding statistic). The test statistic will be:

$$T = \frac{(\bar{X} - \bar{Y})}{S_p \cdot \sqrt{1/n + 1/m}} \tag{2}$$

where $\bar{X}, \bar{Y}$ are the sample means of Canadian-born and immigrant workers' wages, respectively, n is the number of Canadian born workers in the sample, m is the number of immigrant workers in the sample, and $S_p^2 = \frac{(n-1)S_X^2 + (m-1)S_Y^2}{n+m-2}$ is the pooled sample variance (S_X, S_Y are the sample variances). This test statistic follows a t_{n+m-2} distribution. We are going to do a 2-sided test with $\alpha = 0.05$, so our statistic is $|T|$. Simply plugging in numbers into this formula (you could use Excel to get means and standard deviations of data) we obtain a t-statistic 0.8202. We look up our p-value using a T-distribution with 998 degrees of freedom (approximately normal), and get a p-value of 0.4123. Therefore, we fail to reject the null that Canadian-born and immigrant workers earn the same.

(ii) Here is the output for the test in Stata: as expected it is the same as the test computed manually.

Two-sample t test with equal variances

Group	Obs	Mean	Std. Err.	Std. Dev.	[95% Conf.	Interval]
0	233	29093.31	3756.1	57334.38	21692.89	36493.74
1	767	32267.48	1802.369	49916.2	28729.31	35805.65
combined	1000	31527.9	1635.757	51727.17	28317.98	34737.81
diff		-3174.166	3870.029		-10768.49	4420.161

```
    diff = mean(0) - mean(1)                                      t =  -0.8202
Ho: diff = 0                                     degrees of freedom =      998

    Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
 Pr(T < t) = 0.2062         Pr(|T| > |t|) = 0.4123          Pr(T > t) = 0.7938
```

b) To perform a test of whether there is wage discrimination against women, I am going to do a difference in means test on a subset of men and women who I think have similar productivity. The goal here is to control for characteristics that differ between men and women, and affect the productivity of the worker. We say there is discrimination if two equally productive workers are paid differently because of a characteristic that does not affect productivity (ex: gender). So we choose our controls in order to create a sample of workers that we feel are the same in terms of productivity.

First I construct an hourly wage variable, to control for differences in part/full-time work. I also only include men and women who have a positive wage (i.e. are in the work force). My other controls for this test are: canborn==1, age between 18-30, and education is a BA degree. Here is the output:

Two-sample t test with equal variances

Group	Obs	Mean	Std. Err.	Std. Dev.	[95% Conf.	Interval]
0	34	24.14473	3.138683	18.30151	17.75903	30.53043
1	43	22.22115	1.862879	12.21571	18.4617	25.98059
combined	77	23.07052	1.723884	15.12702	19.63711	26.50393
diff		1.923582	3.487565		-5.024003	8.871168

```
    diff = mean(0) - mean(1)                                      t =   0.5516
Ho: diff = 0                                     degrees of freedom =       75

    Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
 Pr(T < t) = 0.7086         Pr(|T| > |t|) = 0.5829          Pr(T > t) = 0.2914
```

Note that I have only chosen one subset of my observations; you should perform several tests using different age groups and education groups to see if you obtain similar results. In this case, we fail to reject the null that male and female wages are equal. Therefore does this mean we have no wage discrimination? The answer is we really can't say. Although we have controlled for observables (whether or not these are the all the relevant observables could be argued), there could be unobservables that cause differences between the productivity of men and women. One example would be ability; if men and women differ in ability (which we do not have a measure of in our data), than that could cause differences in productivity between men and women even after controlling for observables. The point is we cannot get a definitive answer on discrimination using a simple test such as the one above; we need to try make some claims (perhaps using economic theory) about unobservable differences between the productivity of men and women as well.

Stata Commands:

```
//Problem 2//
//part a, ii)
//test that immigrant mean earnings are equal to candian earnings

// variable canborn is 1 if canadian born, 0 otherwise

ttest wages, by(canborn)

//part b
gen wage_hour=wages/(hrswrk*wkswrk)
ttest wage_hour if ba==1 & 18<age<30 & canborn==1, by(female)
```

Problem 3

We want to know whether women are more likely than men earn at least \$100,000. Therefore, this can be framed as a one-sided test, with a null hypothesis of $Pr(Wage \geq \$100K|Gender = M) = Pr(Wage \geq \$100K|Gender = F)$, and the alternative hypothesis of $Pr(Wage \geq \$100K|Gender = M) < Pr(Wage \geq \$100K|Gender = F)$. Here is the output from the test:

```
Two-sample test of proportions                  male_pro: Number of obs =      509
                                              female_pro: Number of obs =      491

    Variable |       Mean   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
    male_pro |   .0392927   .0086118                       .022414    .0561715
  female_pro |     .00611   .0035168                     -.0007828    .0130028
-------------+----------------------------------------------------------------
        diff |   .0331828   .0093022                      .0149508    .0514147
             |  under Ho:   .0094823     3.50   0.000
------------------------------------------------------------------------------
        diff = prop(male_pro) - prop(female_pro)                  z =   3.4995
    Ho: diff = 0

    Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
 Pr(Z < z) = 0.9998         Pr(|Z| < |z|) = 0.0005          Pr(Z > z) = 0.0002
```

There are several things to note here. First, to conduct this test (see do-file) I constructed two variables, one for women and one for men, that took on a value of 1 if the observation had a wage over 100,000. Second, this test assumes equal variances from the two populations. The reason this should be the case is that the variance of a Bernouilli random variable (a variable that is 0 or 1 with probability p) is $p(1-p)$. Under the null, the proportions are equal. Therefore, under the null the variances would be equal. Thus, we should assume equal variances. Again, with large n the test statistic follows a normal distribution. From the output, the p-value for our stated alternative hypothesis above is -0.9998; this means we fail to reject the null. The fact that this p-value is so close to 1 indicates that, under the null, the probability of observing a test statistic at least as extreme as the one we obtained *in the direction of the alternative* is nearly 100%! This is because as it turns out, we actually have evidence going in the opposite direction. That is, we can reject the null in favour of the alternative that men have a higher probability of earning over \$100K. The p-value for this one-sided test is 0.0002; we reject the null at the 10, 5, and 1% level.

Stata Commands:

```
//Problem 3//
//create a variable that is 1 if female and earn over 100k
gen female_pro=0 if female==1
replace female_pro=1 if female==1 & wages>=100000

//same thing for men
gen male_pro=0 if female==0
replace male_pro=1 if female==0 & wages>=100000

//test
prtest male_pro == female_pro
```