



Female Chief Financial Officers and the Glass Ceiling

Executive Summary

Analyzing chief financial officer (CFO) compensation data at public U.S. firms, we found that women are paid significantly less than men for the CFO role, even after accounting for a number of factors known to affect pay. We reached this conclusion by first building a model to predict CFO gender, and found that even after controlling for variables like company size, tenure and board memberships, CFO compensation was predictive of CFO gender. We then built a model to estimate the pay differential between male and female CFOs and found that men earn 16.3% more than women on average, even after accounting for a number of variables including market capitalization. We speculate that reasons for the differential may include gender differences in career path and work histories, but testing these hypotheses lies outside the scope of our study.

The Research Project

One nearly sure-fire way to generate some heated discussion is to bring up the topic of the glass ceiling. Originally coined to refer to the lack of women in management, the term's use has been generalized over the years to encompass a variety of obstacles to women's advancement in the work world. Compensation is one area where women often lag their male peers, and remains a hot-button issue. There is little doubt that women are frequently paid less than men. The real question is whether they are truly paid less because they are women, or whether there are factors other than gender that might account for this differential.

At GMI Ratings, we are well situated to explore this question. One of our predecessor firms, The Corporate Library, created a system to collect data on the gender and compensation levels of all directors and officers at Russell 3000 companies. Another of our predecessor companies, Audit Integrity, developed quantitative research methods to analyze governance and accounting data. We have brought those capabilities together to look for patterns in the data, and evaluate their statistical significance.

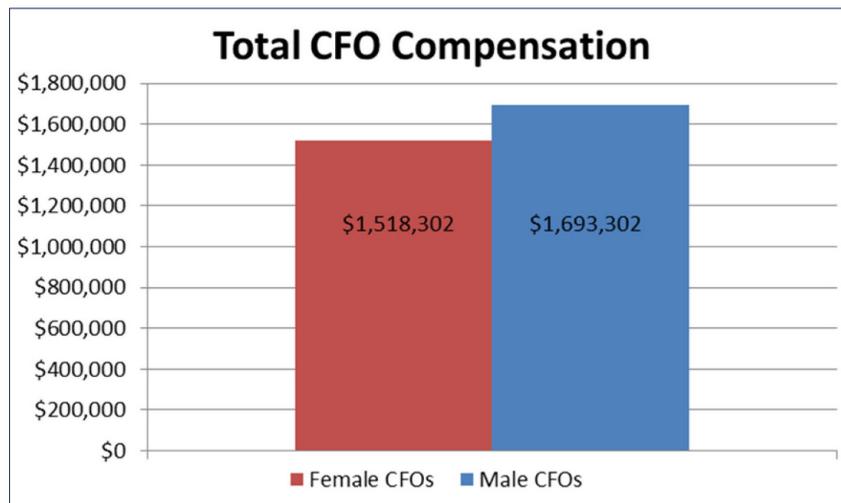
Methodology

To take a fresh look at the question of the glass ceiling, we decided to look at our data on female chief financial officers (CFOs) in the United States and find out whether they were paid the same as their male colleagues for substantially similar jobs. We chose CFOs for two reasons. First, female CFOs represent a large enough sample to be statistically meaningful (about 8% of the companies in our study have female CFOs, while there are far fewer female CEOs). Second, looking at CFOs allows us to focus on a group of people who all have the exact same title, and presumably, have roughly similar job responsibilities and educational backgrounds. The role that gender differences play in determining pay should be simpler to identify in a group that is similar in these other ways.

A conventional approach to answering this type of question would be to look at summary statistics to see if there were statistically significant differences in compensation between male and female CFOs, and then to break down the sample into smaller groups to control for other correlated variables (like company size) to see if the difference in compensation persisted. We decided that an interesting way to consider the question of pay equity would be to look at it in reverse. Rather than looking for differences in pay between women and men, we wondered if we could turn the problem on its head: can we predict people's gender, based on how they are paid? If, in fact, there is no glass ceiling, and all differences in pay between men and women are explained by other variables, then compensation should be of no use in predicting the gender of a CFO. To put this another way: if we set out to build a statistical model to predict gender, testing compensation as well as a number of other variables, and compensation turns out not to be a statistically significant variable in our model once all those other things are included, then we can reasonably conclude that gender pay equality (at least among CFOs) has finally become a reality.

Total CFO Compensation

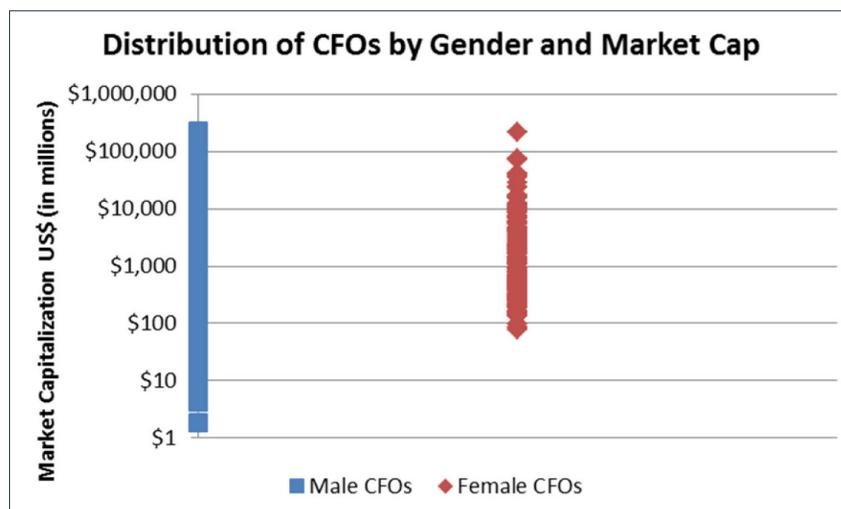
To begin, we first looked at the absolute pay levels of male and female CFOs in our sample and found that—no surprise—the women, on average, made less.



The more than 10% difference certainly seems significant, but by itself may not be evidence of a gender bias if, for example, women tended to work at smaller companies, where pay for all executives is generally lower than at larger firms.

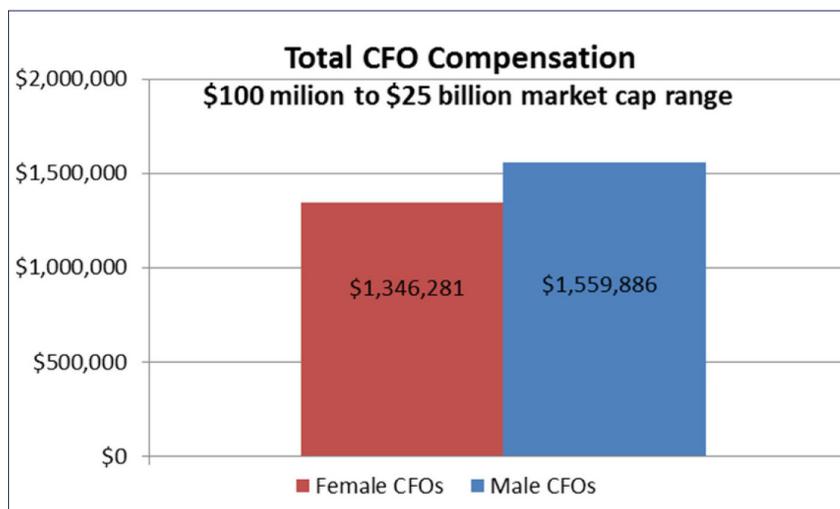
Prevalence of Female CFOs by Market Cap

We therefore, decided to look at the distribution of female CFOs across market-cap groupings, to see if they do indeed tend to serve at smaller companies.



We found that female CFOs are extremely rare at the two extremes of market-cap—there are hardly any women serving as CFOs at companies with under \$100 million or over \$25 billion in market capitalization. Accordingly, we decided to restrict our study to companies between \$100 million and \$25 billion in market cap in order to prevent our results being skewed by the small sample of female CFOs outside of that market cap range.

We then revisited our look at Total CFO Compensation, this time limiting the comparison to companies in our \$100 million to \$25 billion market cap range.



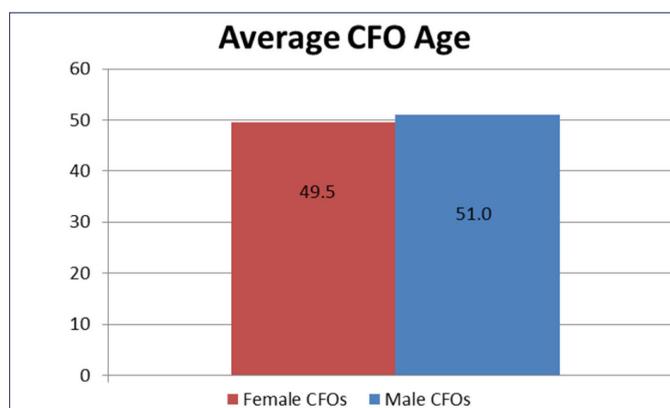
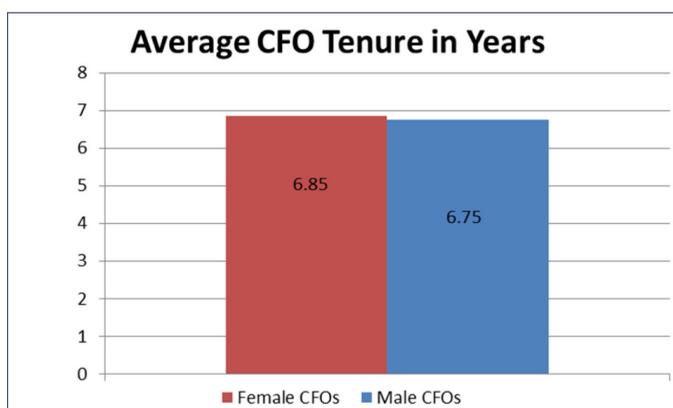
We found that the pay gap between male and female CFOs increased from just over 10% to about 14% when we focused on companies in this slightly narrower market cap range.

Initial Regression Results

In our first attempt at modeling, we tested both market cap and CFO compensation and found that CFO compensation is a marginally significant predictor (at the 10% level) of gender: the more a CFO makes, the less likely that person is to be female. *(The statistical results of this regression are presented in the Appendix.)*

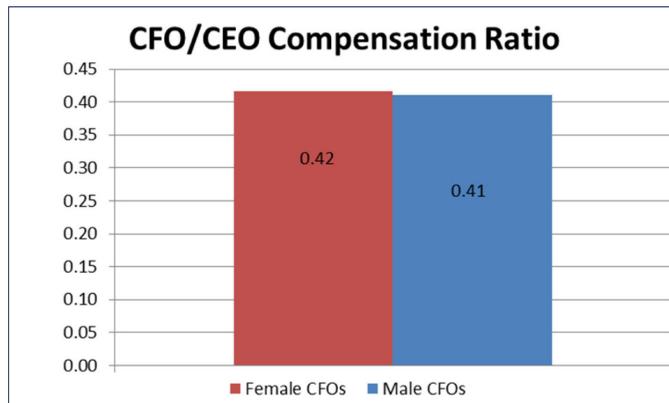
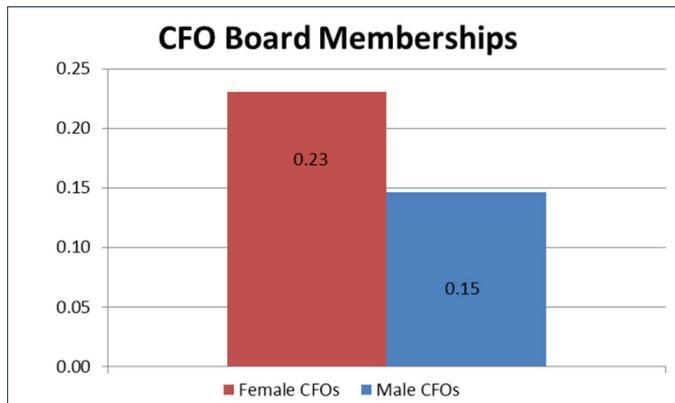
Age and Tenure

Again, by itself, this is not evidence of gender discrimination per se. Perhaps women, as a group, tend to have some other characteristic that is less valued in the marketplace—for example, perhaps they are younger or have shorter tenures in their positions, and therefore earn less. So we set out to look at those variables. As the chart below shows, though, there is virtually no difference in average tenure (in their positions) between the male and female CFOs. Looking at CFO age, however, we find more of a difference—male CFOs are on average about 18 months older than female CFOs.



Additional Variables

We then cast a wider net for other variables that might help explain the pay differential, such as the number of boards CFOs serve on. Interestingly, female CFOs tend to serve on significantly more boards than their male counterparts; however, overall it is not common for CFOs to serve on boards. We then looked at the ratio between CFO and CEO compensation at each company, and identified the median ratios for the group with male CFOs and the group with female CFOs. There was little difference between the two.



Finally, we re-ran our model testing each of the variables we had examined. We found that CFO compensation, CFO age, and number of boards were all statistically significant predictors of CFO gender. Our model indicates that the more a CFO is paid, the less likely the CFO is to be female.

Evaluating the Model

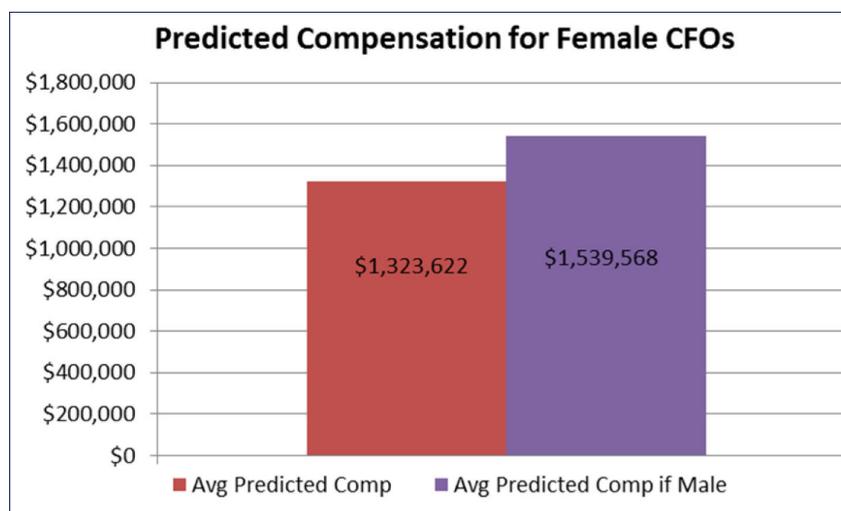
So, we'd built a model to predict CFO gender; then what? Well, one obvious question was: "How good is the model?" To answer that question, we used a test of concordance which looked at every possible combination of one female CFO and one male CFO in our dataset. For each combination, we looked at the probability of being female assigned to each member of the pair by our model. If the female CFO was assigned a higher probability of being female than the male CFO was, then that observation was concordant, or a "win" for the model. If the male CFO had a higher probability of being female, then the observation was discordant, or a "loss" for the model. Sixty-two percent of the time our model correctly predicted which CFO was female. *(Statistical data to support this finding is in the Appendix.)*

Given that any model would randomly be correct 50% of the time, 62% is not necessarily an impressive result. However, we weren't trying to build a perfect model to predict CFO gender--we were simply trying to demonstrate that a model that includes CFO compensation is predictive of CFO gender. If CFO pay predicts CFO gender, then it becomes very difficult to argue that pay equity exists across genders for CFOs.

Estimating the Pay Differential

While technically our 62% success rate in predicting CFO gender may be evidence of the existence of a glass ceiling in pay, it doesn't speak to the size of the pay gap. To address this issue, we created a different kind of model. Rather than predicting CFO gender, this time we attempted to predict CFO compensation using CFO gender as one of the input variables. With this simple linear regression approach, the coefficient we'd get for the CFO gender variable would essentially put a price tag on what it means to be a female CFO vs. a male CFO.

The details are shown in our Appendix, where the Parameter Estimate column is the coefficient for each variable and indicates that all other things equal (in this case market cap, number of board memberships, age and CEO compensation) the average female CFO would earn about \$215,000 more if she were male. This number was arrived at by changing only the gender variable in the model for female CFO pay. It therefore shows how women would fare if all their other characteristics, such as age and number of boards, remained the same, and only their gender were switched.



This would represent a 16.3% increase in average compensation. Looking at each observation individually, we find that our model indicates that the median female CFO would see a 25.5% increase in pay if she were male. The R-squared of this model is .60, indicating that it explains about 60% of the variance in CFO compensation in our sample.

Conclusions

We urge caution in interpreting these results for several reasons. First, while we believe this methodology is a useful way for looking at this issue, our analysis is constrained by the data. Not only do we have a relatively small sample of data, but our sample comes from a single point in time, which may not be representative of a longer period. In other words, perhaps there was something unique about mid-2011 which is exaggerating the pay differential between male and female CFOs. However, this cuts both ways; it is also possible that by chance our dataset is minimizing the pay differential between the genders when compared to historical norms. That said, we do find it striking that female CFOs appear to earn 16.3% less, on average, than their male colleagues after accounting for obvious pay-related factors like tenure, age, and market capitalization of the company.

We speculate that some other factors related to women’s typical work histories may account for some of the pay differential. (Unfortunately we do not have data we could use to test these hypotheses.) For example, our finding that male and female CFOs have similar tenures only relates to their tenures in the position they currently hold. It is possible that female CFOs have shorter work histories, even when they are the same age as men, because they are more likely than men to interrupt their careers to bear and care for children. Similarly, we wonder whether women might differ from men in the routes they typically take as they move up in the business world. For example, it is possible that women are more likely to move up within a single organization over time, while men are more likely to switch employers when they move to the next level of responsibility. Since many companies tend to pay more for new hires than for promotions from within, if men’s and women’s career paths differed in this way, it could explain some of the compensation differential as well.

Appendix

This study is based on post-proxy season 2011 compensation data, pulled in mid-June of that year. In our dataset, we have data on 1,910 Russell 3000 companies (155 of which had a female CFO) with data on total compensation for the CFO and CEO, the age of the CFO, the gender of the CFO, the number of boards the CFO is on, the tenure of the CFO and market cap. We have excluded any Russell 3000 company where either the CFO or CEO had been in the position for less than 1 year, as well as any companies for which we were unable to obtain market capitalization information.

Statistical results of the models referred to in the text are presented below.

Initial Regression Results:

Analysis of Maximum Likelihood Estimates

Parameter	DF	Estimate	Standard Error	Wald Chi-Square	Pr > ChiSq
Intercept	1	-2.2154	0.1383	246.4637	<.0001
cfo_comp	1	-1.39E-7	70913E-8	3.0855	0.0790

Final Model to Predict CFO Gender:

Analysis of Maximum Likelihood Estimates

Parameter	DF	Estimate	Standard Error	Wald Chi-Square	Pr > ChiSq	Standardized Estimate
Intercept	1	-0.1338	0.6850	0.0381	0.8452	
cfo_comp	1	-1.95E-7	8.417E-8	5.3581	0.0206	-0.1471
dirage	1	-0.0426	0.0138	9.5529	0.0020	-0.1562
num_boards	1	0.7208	0.1935	13.8812	0.0002	0.1638

Model Evaluation:

Concordance Test

Association of Predicted Probabilities and Observed Responses

Percent Concordant	62.0	Somers' D	0.258
Percent Discordant	36.2	Gamma	0.262
Percent Tied	1.8	Tau-a	0.039
Pairs	228943	c	0.629

Estimating the Pay Differential:

Parameter Estimates

Variable	Label	DF	Parameter Estimate	Standard Error	t value	Pr > t
Intercept	Intercept	1	-1267378	205425	-6.17	<.0001
is_cfo_female	is_cfo_female	1	-215946	76079	-2.84	0.0046
num_boards	num_boards	1	376170	52451	7.17	<.0001
ln_market_cap		1	225111	19195	11.73	<.0001
ceo_comp	ceo_comp	1	0.17809	0.00572	31.15	<.0001
dirage	dirage	1	7834.36682	3203.98342	2.45	0.0146



www.gmiratings.com