University of Pennsylvania
ScholarlyCommons

Wharton Research Scholars Journal | Wharton School

4-1-2005

Using Simulation to Determine Optimal Employee Participation in Employee Stock Purchase Plans

Alexander Janofsky
University of Pennsylvania

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/wharton_research_scholars/26
For more information, please contact repository@pobox.upenn.edu.

Using simulation to determine optimal employee participation in Employee Stock Purchase Plans

By

Alexander Janofsky
Wharton Research Scholar

Advisor

Professor Brigitte Madrian
Department of Business and Public Policy
The Wharton School, University of Pennsylvania

April 2005

Introduction

An Employee Stock Purchase Plan (ESPP) is a specific type of company stock ownership plan sponsored by an employee's company and defined by appropriate federal regulators. The ESPP is designed to encourage saving and investment among mostly non-executive workers while leveraging the perceived incentive benefits provided by company stock ownership.

ESPPs are a particularly compelling retirement vehicle to study because the recent corporate scandals and subsequent loss of retirement assets have put renewed focus on the allocation of retirement assets to company stock. Many have been calling for government to limit retirement plans which encourage or require investments in company stock. For instance, federal legislators have recently proposed a law allowing investors in 401K plans with assets in company stock to trade out of that stock after three years.

This investigation seeks to explore the risk-return tradeoffs faced by ESPP participants and, in doing so, determine optimal investment behavior.I seek to understand to extent to which the specific benefits of ESPP participation justify concentration of retirement assets into a single stock. Rather than analyze this vehicle from a pure economic perspective to understand how it allocates costs between employees and employers, I focus solely on the employee's investment decisions.

This paper first describes the features, mechanics, and tax treatment of ESPPs. Secondly, after motivating my investigation, I describe the methodology of the simulation I executed to mimic the returns on ESPPs and other assets of interest. Thirdly, I analyze the results from the simulation, including specific comparisons aimed at

determining optimal investment behavior. Finally, I conclude with a suggestion for optimal employee investment behavior in ESPPs.

Description of Employee Stock Purchase Plans

Participation in an ESPP is voluntary and involves after-tax income. Contributed funds are used to purchase company stock at market price or a discount (usually the maximum 15% permitted by regulations). This discount provides ESPPs an instant guaranteed initial return. Thus, failure to participate altogether would be akin to refusing cash. Most plans limit employee contributions to between 10 and 15 percent of compensation.

Generally, ESPPs are open to employee contributions over a specified offering period. The lengths of these periods vary, but many companies choose 6 months. Cash contributions are accrued throughout the offering period directly from salary.

Most of these plans have “look-back” features which specify that the price at which the stock is purchased is the minimum of the fair market value at the beginning and end of the offering period. Some ESPPs also offer employees the option of removing contributed funds from the plan before the end of the offering period.

The look-back feature is an imbedded, at-the-money European call option position with maturity equal to the length of the offering period. Intuitively this follows since the look-back feature only has value if the stock price increases during the offering period. The number of shares on which this option is written is equal to the initial investment (after the discount) divided by the price of the stock at the beginning of the

offering period. Below I derive this result. *T* is initial investment, *d* is the discount percentage, S_i is initial stock price, and S_f is stock price at the end of the offering period.

$$\text{ESPP value at time i} = \frac{\text{T}}{\text{1-d}}$$

$$\text{ESPP value at time f} = \frac{\text{T}}{\text{1-d}} Max\{\frac{S_f}{S_i},1\}$$

$$\text{Value of the lookback feature} = \frac{\text{T}}{\text{1-d}} Max\{\frac{S_f}{S_i},1\} - \frac{\text{T}}{\text{1-d}}$$

$$= \frac{\text{T}}{\text{1-d}}\left(Max\{\frac{S_f}{S_i},1\} - 1 \right) = \frac{\text{T}}{\text{1-d}} Max\{\frac{S_f}{S_i} - 1, 0\}$$

$$= \frac{\text{T}}{\text{1-d}} \frac{1}{S_i} Max\{S_f - S_i, 0\}$$

Clearly this payoff represents a European call option on the number of shares that could be bought with the discount at the beginning of the offering period, with strike equal to the initial stock price.

For tax treatment, ESPPs are classified as either qualified or nonqualified. Qualified plans comply with specific tax rules[i]. Among these rules is an annual contribution limit of $25,000 for each employee and the 1-2 holding rule[ii] which ensures that participants have held onto the stock for a sufficiently long period of time.
All gains on ESPP investments that are sold immediately after the purchase occurs at the end of the offering period are taxed as income. Qualified plans are taxed so that gains from the discount are taxed as income, but returns in excess of the discount are taxed as capital gains. For instance, if the sale position is less than the original position after the discount but still includes a gain, the entire gain is taxed as ordinary income. Since the capital gains rate is usually below the marginal income tax rate for ESPP investors, a qualified position provides a tax advantage. Nonqualified plans are taxed similarly to

qualified plans except the discount and benefits from the look-back are both taxed as ordinary income regardless of if, at disposal, an overall gain or loss on the ESPP is realized.

If one considers that the discount is just ordinary compensation, the tax benefit from an ESPP investment over an independent investment in company stock outside of the ESPP framework ultimately comes from the fact that benefits associated with the look-back option are taxed at the capital gains rate.

ESPPs differ from more ubiquitous 401(K) plans in several ways. While ESPP contributions come from after-tax income, 401(K) contributions are pre-tax. Furthermore, ESPPs allocate all contributions to company stock while 401(K) plans may include some company stock as one of numerous investment options. The ESPP look-back and discount features are also not part of a standard 401(K) plan. However, 401(K) plans often contain employer matching provisions which are not common in ESPPs. The two vehicles also substantially differ in tax treatment.

Question of Interest

The primary disadvantage of ESPP participation is it concentrates a large amount of personal investment capital in a single risky security, undermining portfolio diversification. The benefits of diversification are well-established. They stem from the fact that a portfolio consisting of several uncorrelated return streams has a lower variance in total return than a portfolio with the same expected return but consisting of only one or two return streams. Since many investors are sufficiently diversified, the market prices equity to reflect the marginal risk-reward tolerance of these diversified investors. Thus,

underdiversification undermines the relative value of a portfolio by exposing the investor to the risk of an undiversified investment but yielding the lower return demanded by diversified investors. So, concentrating personal investment capital in a single stock undermines the value of that investment in terms of the Capital Asset Pricing Model.

My intent in executing the simulation described below is to generate distributions of returns from ESPP investment and as well as from some standard alternatives. By comparing the generated return distributions, I will be able to describe the specific nature of the risk-reward tradeoff of ESPPs as compared to alternative strategies and benchmarks. I hope to understand how this tradeoff evolves over various time horizons. Since participation in ESPPs clearly is advantageous in absolute terms, I am most interested in discerning optimal participation length. For instance, is same-day sale the optimal strategy or should participants maintain their ESPP position for longer periods.

Another less-obvious economic disadvantage to investing in one's own company is the correlation between labor income and the company stock price. I do not address this because the simulation focuses on returns to ESPP investment while excluding a more robust life-cycle approach. Analysis of the labor income correlation effect would require imbedding the simulated ESPP investment in a life-cycle model to track labor income, investment income, and total wealth over a long period of time.

Methodology

To determine the employee's optimal participation behavior with regard to Employee Stock Purchase Plans, I simulate returns to participation over varying investment period lengths for the entire stock market. This simulation specifically

addresses the following question: assuming a randomly selected firm implemented an ESPP, what characteristics would define the resulting return distribution for participating employees, and how would this distribution compare to distributions of benchmark returns?

I choose to examine all firms instead of limiting my analysis to those with a record of offering ESPPs for two reasons. Practically, I would have found it difficult obtaining a list of companies that offer actual ESPPs and then obtain the specific parameters and offer dates of those plans. Furthermore, I believe that characterizing the risks and rewards from the perspective of employees for a company interested in implementing an ESPP will yield more fruitful conclusions.

The one drawback to simulating over all firms as opposed to just those which actually offer ESPPs, is such a method ignores potential biases in the type of firms which choose to offer ESPPs. For instance, perhaps firms with a certain type of return distribution are more likely to offer ESPPs. Since I am focusing on the larger question of how an employee in any given firm would respond to an ESPP offer, this bias should not affect my conclusions.

Raw Data

The only raw data necessary for these simulations is a list of all publicly traded companies, their stocks' monthly holding period returns, and the simultaneous returns on the Standard & Poor's 500 Index and on a riskless investment. I choose to examine holding period returns because price is affected by splits, does not account for mergers, and ignores dividends. Monthly data is most appropriate because the parameters of ESPPs dictate a minimum investment period of six months. This means yearly data

would not provide enough detail while daily data would incur great computational expense without a substantial increase in the simulation's validity or accuracy.

The list of all publicly traded companies is easily obtained from the joint COMPUSTAT/CRSP database. Using the joint database is important because, while CRSP provides all of the necessary stock return information, it its indexes records by security instead of by company. Unfortunately, the joint database only extends back to 1962 instead of 1925, as the CRSP database proper does. As a result, my analysis of "all companies" is limited to companies which have been publicly traded at any time from 1962 to 2004 and which have been assigned a PERMNO identifier. The resulting list contains 8,487 companies. Since I only use the joint database to retrieve company identifiers, if a company also has return data prior to 1962, I include that in the simulation. For example, IBM is included in my list of companies since it has been traded since 1962, but my IBM monthly return series extends back to 1925.

While, ideally, this analysis would have included every single publicly traded company over the past 85 years, the sample universe is still extremely large and I have no reason to believe that any of the aforementioned exclusions bias results. Specifically, no evidence exists of a relationship between returns and whether or not a company has been assigned a PERMNO. Additionally, forty years will likely be long enough to account for business cycle bias and survivor bias. Business cycle bias comes from looking at returns from a shorter period of time and thus characterized by a particular macroeconomic environment. For instance, limiting the simulation to 1995-1999 would probably characterize an ESPP investment as more rewarding than it perhaps would be in a more neutral macroeconomic environment because equity returns during the late 90's were at

an historical peak. Additionally, limiting the simulation to a shorter, more recent series of returns would implicitly bias the results toward “survivor” companies. Companies which had failed before the data series began would be not be included, while companies which had succeeded would be included. A longer data period helps control for both business cycle and survivor bias.

I realize that there is a tradeoff between using older data to control for business cycle bias and survivor bias in that the equity markets have probably undergone significant structural changes over the past 40 years. If the relationship between individual equity, market, and fixed income returns has changed over the last 40 years, the predictive value of my analysis would be undermined since I would be using properties that have since changed to describe the current and future world. An example would be the theoretical premium on equity returns over the risk free rate. Since the preponderance of evidence suggests that this premium has decreased during the past 40 years, the spread between equity returns and a riskless return might be significantly higher in the simulation than in the current or future market. Other structural changes might include increased liquidity and returns distributions beginning to closely resemble Gaussian white noise (i.e. reduced autoregression). I reconcile the simulation’s vulnerability to bias from structural changes in the market by positing that the effects of survivor and business cycle bias are likely much larger and more detrimental to my specific analysis than the effects of structural biases.

The decision to include data prior to 1962 where available is less clear. I feel that this decision will not affect my results because the number of companies affected is likely small compared to the size of the sample population. Additionally, this does not create

any obvious bias. Time and computational constraints prevent me from performing the simulation a second time on a data set which does not include any data prior to 1962.

The return on the S&P 500 is a simple and ubiquitous proxy for the market return. This is a useful benchmark because it represents equity returns given full diversification. Since lack of diversification is one of the main disadvantages to ESPP participation, such a benchmark will help isolate this effect.

The risk free rate for a given month-long period is estimated by the Fama risk free rate, a CRSP data set which utilizes appropriate historic t-bill rates. It is useful for simulating a baseline return to control for the effect of nominal interest rates on investment allocation decisions.

The resulting raw data for each company consists of a sequence of 3-tupples whose entries are monthly returns to company stock, the S&P 500, and a riskless investment.

Simulation Description

The simulation generates stock return series for each company by randomly drawing (with replacement) a 3-tupple to represent each simulated month. The sample space for each step is in bijection with the entire sequence of monthly returns for the company. Drawing by 3-tupple instead of separately for each asset maintains the historic correlations between the three assets. This simulation procedure is, in part, derived from a similar procedure used by Brown, Liang, and Wiessbenner (2004)[iii].

This method implicitly assumes that future returns are distributed identically to past returns. While the veracity of such an assumption is doubtful, it seems no less reasonable than any other assumption regarding the predictive value of statistical

analysis. Non-stationarity, a virtual certainty over 40 years, would clearly undermine my conclusions since it would invalidate the means and variances of the simulated distributions. However, since I am more concerned with the relative properties of the distributions as opposed to their absolute predictive power, I believe that this method is reasonably valid.

The simulation simulates 60 months (5 years) of returns on the three assets. I chose this length of time because, on a long enough horizon, total returns to an ESPP investment resemble a standard investment in company stock. I wanted to avoid expending computational energy generating data less central to the goal of evaluating ESPP participation. I chose not to simulate a shorter period because I wanted to run the process long enough for the ESPP to "converge" to the standard company stock investment. To ensure that the sample space from which the simulation draws is large enough, I remove companies which do not have at least 60 months of data. While the exact cutoff is somewhat arbitrary, the reasoning is clear in the extreme: a company with only 2 months of data would certainly pollute my analysis by putting a tremendous amount of weight on two months of returns which would individually not represent the theoretical return distribution to a useful extent. On the other hand, I must be sure not to eliminate too many companies or even to introduce additional survivor bias. The final population list contains 6,315 companies to which I apply the simulation.

I designed the simulation to capture the most prevalent incarnations of the ESPP. To that end, I assume the following: an offer period of 6 months, a standard look-back feature, a 15% discount, qualified disposition after the 18 months from the start of the

offer period, contributions only made on the first day of the offer period. As noted above, these features would describe the plurality of ESPPs very accurately.

I have decided to only assume one contribution because I am mainly interested in optimal employee behavior with regard to any single ESPP offering. Additionally, I could have used the simulated return distributions to replicate returns from further contributions if I later deemed it necessary. Finally, I can perform the single investment simulations without assuming a specific monetary amount for wealth and contribution. Transcending such parameters contributes to the robustness of this analysis.

Total return for each asset is merely the gross monthly return times the gross total return prior to that month. Dividend reinvestment is assumed. Throughout this analysis I focus on returns instead of absolute dollar amounts and prices because all of the information I use and all of the processes I simulate are much more efficiently represented by and executed in terms of returns. Additionally, focusing on returns allows me to avoid having to assume an initial investment amount, thereby making the analysis more robust (as previously noted).

I ignore transaction costs because the shortest trading period I consider is six months and because no investment position I simulate involves more than three trades. Given current transaction costs for equity investments and the horizon that interests me, the effect of transaction costs should be negligible if asset allocation is not rebalanced. Moreover, transaction costs should effect all equity investments equally. At worst, spreads with the riskfree rate would be affected.

I ignore inflation because this analysis is comparative rather than absolute and inflation would affect all asset returns relatively equally.

Simulation Algorithm

The following describes a single simulation run for a given company. Each simulation run generates a vector containing after-tax total returns for various assets and holding period lengths.

At the contribution date, investors begin with a gross return of 100% on their initial investment. ESPP investors immediately receive a return equal to 1/(1-discount) to account for effect of the discount on their total return. Here I assume a 15% discount as is most common among ESPPs offered.

Each month, a tupple containing monthly returns for the three assets is randomly drawn from the space of such tupples for the given company. For each asset, the gross monthly return is multiplied times the existing gross total return and then this step is repeated for as long as specified.

To simulate the look-back feature for ESPP investors, after 6 months of returns have been generated, if the total ESPP return is less than the initial return from the discount, ESPP return is reset to the initial return from the discount. This is the total return equivalent to purchasing the company stock on the last day of the offer period given that the stock price fell over the course of the offer period.

Total returns to investment in each of the three assets are reported after 6, 18, 36, and 60 months from the contribution. I chose these investment lengths because six months is the first date investors can liquidate their ESPP positions and still benefit from the discount and the look-back feature, the two aspects of ESPPs that are clearly advantageous regardless of the investor's preferences or situation (i.e. refusing them would be turning down pure compensation). Eighteen months is the first date where the

ESPP position is considered a qualified disposition and therefore eligible for more favorable tax treatment. Sixty months is seemingly long enough to demonstrate the effects of holding an underdiversified position and thirty six months is a reasonable observation point between 18 and 60. Again, computational constraints prevent me from observing more data points.

For each asset class, appropriate tax treatment is applied to the total returns for each of the four different simulated holding periods. To preserve simplicity without undermining the value of this analysis, I assume a marginal income tax rate of 28% and a capital gains tax rate of 20%. Since all analysis is comparative, the absolute tax rates matter less then the spread between the two. Given the likely income characteristics of ESPP investors and current tax policy, both the 8% spread and the absolute rates are realistic assumptions. As described above, gains on nonqualified ESPP positions are taxed at the income tax rate while gains on qualified ESPP positions beyond gains from the discount (1/(1-discount) – 1) are taxed at the capital gains rate. Investments in the market portfolio (S&P 500) are taxed at the capital gains rate to conform to reality. Since the risk free investment is intended to simulate return from a certificate of deposit or a money market account, it is accordingly taxed as income.

I chose to ignore the tax benefits associated with capital gains losses. In part this is realistic because the types of individuals who would participate in ESPPs are unlikely to have gains needing to be offset since they probably do not much capital tied up in other investments. Additionally, given the goals and scope of this simulation, evaluating the implications of such benefits would have been ambiguous and distracting.

After taxes are taken into consideration, the total returns for each of the four holding periods are complete. Next, a portfolio of three assets is generated according to a 1/n allocation without rebalancing. The intent of this portfolio is to provide the hypothetical investor with an option which captures some of the benefits of ESPP participation with significantly less risk. Since n=3, one third of an investor's hypothetical assets are allocated to each of the three assets. Because I ignore rebalancing, the investor earns the average of the three after-tax total returns for each period. The no rebalancing assumption is legitimate since previous studies overwhelmingly indicate that middle class investors rarely have the time, knowledge, or accommodative transaction costs to regularly rebalance their portfolios. While it is possible that such an investor might tweak his allocation over the course of five years, I choose to ignore this for simplicity and because I believe the non-rebalanced portfolio does not significantly depart from reality. The 1/n assumption is less valid but studies do suggest that such an allocation roughly approximates the behavior of many individual investors.

Other Considerations

Because of the sheer size of the data set, I am unable to fully explore the optimal simulation length. I decide to perform the simulation 10,000 times for each of the 6,315 companies in the data set because this number of trials severely tested the realms of feasibility given my computing constraints and yet still satisfied reasonable convergence requirements. Convergence is not a primary concern because the aggregate data I analyze is likely to converge much faster than the pure stock return distributions for a particular company.

Results and Analysis

Having generated 10,000 sets of total return data for each of the 6,315 companies, I must now measure and analyze aggregate characteristics of the distributions of returns.

Percentiles

To acquire a coarse initial description of the return distributions, I compile a table containing the mean percentiles of the gross returns **(Table 1)**. Each entry in the table is calculated by sorting the 10,000 simulated return entries for each asset and holding period of each of the 6,315 companies. The appropriate percentiles are then identified for each asset and holding period of each company. Finally, the percentiles for each asset and holding period are averaged over the 6,315 companies. In this way, I characterize the aggregate distributions of gross returns. A table entry is thus a proxy for the probability that the gross return on a given asset and holding period will be less than the entry.

Table 1 also contains the means of the means and standard deviations for each asset and holding period. With these observations, I can analyze the skew and variability of the return distributions. **Table 2** merely translates **Table 1** into annualized net returns with monthly compounding.

Analysis of these two tables suggests that the data contain most of the properties I assumed the data would display. For instance, the median (50th percentile) annualized net return of the S&P 500 stays relatively constant over all the holding periods. This is as expected since, having controlled for business cycle risk, no evidence suggests that the average yearly return on the S&P 500 should change, regardless of the time horizon used. Also as expected, the standard deviation of gross returns of all three assets increases as holding period increases. This is consistent with intuition because the likelihood of

extreme total returns, such as the loss of the entire initial investment or a tenfold increase in the initial investment, increases dramatically over time. Moreover, note that 6 month returns to the ESPP are bounded from below in both tables. This is because of the imbedded look-back feature. According to the distribution, this feature is in-the-money between 25% and 50% of the time. Such a result it suggests that, over 6 months, any given stock is likely to increase in value between 25% and 50% of the time, reasonable aggregate behavior given generally accepted knowledge of equity returns. The median annualized net return for the ESPP investment decreases as holding period increases because a longer holding period decreases the influence of the lower bound provided by the look-back feature.

The most important general characteristic of the percentile tables is the fact that the mean returns on the ESPP and the portfolio investments are significantly higher than the mean returns on the comparable S&P 500 investments. On the other hand, the returns on the ESPP and the portfolio are significantly more variable than the returns on the benchmark. Thus, the data, at least roughly, illustrates the risk-return tradeoff of concentrating a significant portion of one's investment in company stock as opposed to a diversified market portfolio.

The ESPP gross return distributions have extremely long and wide right tails. This is explained by two properties of the ESPP returns. The first is the influence of the look-back feature. The second is the fact that loss potential is limited while gain potential is unlimited. While gross return cannot dip below 0, there is no upper limit. For instance, some of the simulated paths even resulted in ESPP returns of 6000% over a period of 60 months.

While these tables are useful for gleaning general properties of the return distributions, they do not provide sufficient detail to enable deeper comparisons. Such specific comparisons require more refined estimates for the aggregate return distributions and are necessary to inform optimal ESPP participation behavior.

Densities

In order to obtain visual representations of the return distributions, I utilize nonparametric kernel density estimation. I use a Gaussian kernel function and estimate 200 data points between 0 gross return and 12 gross return. I choose 12 as the upper bound because at least 95% of the area of all 12 return distributions is contained within the 0 to 12 interval. A benefit of using the same exact interval and data points for each of the density estimations is it enables easy aggregation and comparison.

For each company, I use the density estimation procedure to estimate the 200 data points for each of the 12 return distributions. For each of the 12 return distributions, I then average the 200 data points over the 6315 companies. The result is an aggregate density for each asset and holding period pair. Next, I employ a spline procedure on the 12 aggregate densities to obtain a piecewise analytical formula for each. Finally, I normalize each analytical formula by its 0 to 12 integral to ensure that the area underneath is identically 1.

While this procedure requires some unorthodox assumptions, it seems a reasonable method for obtaining aggregate distributions of returns. Specifically, visual representations of these distributions provide a complete picture of the benefits and uncertainty associated with the various investment strategies considered.

Figures 1-4 plot the three gross return distributions for each of the four periods. **Figure 1** shows the 6 month returns. As expected, the ESPP investment seems to provide the highest returns with a limited downside because of the look-back. The density distribution is not exactly vertical at gross return 1.12708 because the kernel density procedure divided the weight concentrated at that point between two nearby points. The S&P return seems to be distributed roughly lognormal, as expected. It also appears to have a smaller variance than the ESPP. Moreover, the portfolio return distribution exhibits characteristics of both pure return distributions.

Figure 2 shows the 18 month returns. Already the downside risk eliminated by the look-back expresses itself. Clearly, the ESPP return distribution is much wider than the return distributions of the other two assets. While the ESPP has, by far, the highest upside, it also is the only asset that offers a realistic chance of 0 gross return (a total loss of investment). Note that, although it contains the extremely risky ESPP investment, the portfolio appears to have even less downside risk than the S&P. Both the portfolio and ESPP still appear to be skewed because of the residual benefits of the discount and look-back.

Figure 3 shows the 36 month returns. The S&P return is almost symmetric by this point while the ESPP return continues to widen dramatically. Intuition suggests that the ESPP return should resemble a Gaussian distribution truncated at zero to reflect investments that collect at zero, a point which that cannot go beyond or recover from. The kernel density procedure, however, results in flatter left tail. Again, the portfolio appears to have less downside risk than the S&P while offering a significantly greater upside.

Figure 4 shows the 60 month returns. The distortion caused by the smoothing procedure is even more pronounced in the density estimation of the 60 month gross returns. Again, based on intuition, the plot should resemble a normal distribution truncated at 0 with a very high density at 0 itself. Instead, it seems that much of the density within epsilon of 0 has been spread to other points, creating the illusion of a bimodal distribution. The S&P continues to converge to the normal distribution, as does the portfolio to some extent. The portfolio still maintains its long right tail, reflective of extreme values in the ESPP return distribution.

While visual approximations of return distribution densities are informative and interesting, distinguishing between the economic characteristics of the relevant investments requires the introduction of a utility framework and an assumption about preferences

CRRA Utility and Scenario Comparisons

In order to inform optimal behavior, I must first assume preferences under which employees can maximize utility. I choose to limit my focus to the Constant Relative Risk Aversion preference structure.

$$W = \text{total wealth}$$
$$c = \text{coefficient of relative risk aversion}$$
$$U(W;c) = \frac{W^{1-C}}{1-c}$$

This CRRA utility function exhibits the unique property that the coefficient of relative risk aversion reflections the risk aversion of the employee. For instance, when c=0, U is linear so the employee is risk-neutral. As c increases, the employee becomes increasingly risk averse. As c approaches 1, U converges to log utility.

Because of this property, rather than erroneously comparing absolute utility levels, I instead compare return distributions by solving for the coefficient which equates the expected utilities that result from each investment strategy. Thus, for a given pair of return distributions, if the c which equates expected utilities is greater than the level of risk aversion for a given employee, the riskier strategy is preferred. This is because the employee would have to be more risk averse than he already is to prefer the less risky strategy with its inferior expected return to the riskier strategy with its superior expected return.

Since the spline procedure produces unwieldy piecewise analytical functions, I cannot apply the CRRA analysis directly to the generated density estimates. Instead, I use the generated densities to create inverse cumulative distribution functions and these to randomly generate 100 points which represent each distribution. Next, for each distribution, I assume an initial investment and other wealth at the end of the holding period to generate total wealth given each gross return. Note that the absolute levels of initial investment and other wealth are irrelevant. Only the relative size of initial investment and other wealth matter. For instance, regardless of the dollar amounts assumed, the same c results from the initial investment, other wealth pairs of ($3,$10) and ($3000, $100,000). This allows me to generalize my findings to the relationship between initial investment and other wealth.

Finally, since this procedure only allows for binary comparisons, I must decide which investment strategy pairs to use. Since I am primarily interested in discovering optimal holding period length, I compare investment strategies of equal total holding period but whose ESPP exposure increases according to the time increments I simulated.

I begin with a 60 month investment in one of the two benchmarks. I then compare this to a 60 month strategy consisting of the first 6 months in the ESPP and the remaining 54 months in the benchmark. Then I compare the 6 month ESPP strategy to a similar strategy with the first 18 months in the ESPP followed by 42 months in the benchmark. This continues until I compare a 36 month investment in the ESPP to a 60 month investment in the ESPP. I assume a constant total holding period because comparing returns over unequal periods would bias results due to the time value of money and positive expected returns for all investments involved. For instance, all else being equal, a 36 month investment would yield a lower expected return than a 60 month investment.

I perform the aforementioned analysis on two scenarios. The scenarios correspond to the S&P benchmark and the portfolio benchmark. Subsequently, I analyze the differences between the two scenarios. In order to generate meaningful conclusions, I assume employees have a CRRA coefficient of 3. This assumption is arbitrary but reasonable. Obviously, applying this analysis to "real world" situations would require a more careful coefficient assumption.

Table 3 contains the results for Scenario 1, in which the S&P investment is the benchmark. I have assumed an initial investment of $3000 and consider different other wealth assumptions, between $0 and $300,000, in each column. Each row represents a pair of investment strategies being compared using the CRRA procedure. The strategy pairs increase in variance down the rows.

Interestingly, the coefficients in each column decrease down the rows. This implies that the incremental differences between the less risky strategies are smaller than the incremental differences between the riskier strategies. Thus, the increased return

sacrificed by choosing the 6 month exposure over the 18 month exposure is greater than the return sacrificed in choosing the 36 month exposure over the 60 month exposure. This is particularly interesting because, in terms of number of months of exposure, the riskier comparison clearly forgoes more months of high return exposure than the low risk comparison.

More predictably, the coefficients increase as other wealth increases. This means that individuals of greater wealth need to be increasingly risk averse to prefer the less risky strategies. Such behavior results because large reserves of other wealth essentially constitute diversification. Thus, as other wealth increases the marginal increase in the variability of total 60 month wealth due to ESPP investment decreases.

The first row of **Table 3** compares the 6 month ESPP followed by 54 month S&P strategy to the 60 month S&P strategy. For all wealth assumptions, the CRRA equating these two returns is large (and most likely does not exist). This is because the 6 month ESPP strategy offers less downside risk than the S&P strategy, because of the look-back, and a much higher expected return, due to the discount and the fact that single equities yield higher expected returns than the market portfolio. The next row compares the 18 month ESPP strategy to the 6 month ESPP strategy. While investors with $0-$10,000 in other wealth after 60 months would clearly prefer the less risky 6 month exposure, those with more than $50,000 in other wealth would prefer the riskier 18 month exposure. The third row compares the 18 month ESPP exposure to the 36 month exposure. Here again, those with less than $10,000 in other wealth prefer the less risky, 18 month, strategy while those with more than $50,000 in other wealth prefer the riskier, 36 month, strategy.

The final row compares the 36 month exposure to a full 60 month exposure. In this case, those with less than $50,000 in other wealth prefer the less risky 36 month strategy.

Thus, given my assumption of c=3, I can roughly conclude that those who expect their other wealth, after 60 month, to be less than 4 times their initial investment should only invest in the ESPP for 6 months and then move their money to the more diversified market portfolio. Similarly, those who expect their other wealth to be less than 17 times their initial investment should not invest in the ESPP for more than 36 months. Those who expect their other wealth to exceed 30 times their initial investment should invest in the ESPP for at least 60 months if not longer.

Table 4 contains the results for Scenario 2, where the portfolio is the benchmark instead of the S&P. Other than the benchmark variation, the structure and assumptions of **Table 4** are exactly the same as those of **Table 3**. As in the first scenario, the coefficients decrease down the rows. Also, the coefficients increase as other wealth increases.

The optimal behavior implied the c=3 assumption is similar. Those who expect their other wealth to be roughly 10 times or less the amount of their initial investment should only invest in the ESPP for the minimum 6 months. Those who expect their other wealth to be roughly 60 times or less the amount of their initial investment should not invest in the ESPP for more than 36 months. Those who expect to have even higher other wealth can probably invest in the ESPP for longer, if not indefinitely, and still maximize expected utility.

Conclusion

This paper described a simulation procedure designed to uncover optimal employee behavior given the choice of participating in an Employee Stock Purchase Plan. An ESPP is a savings vehicle which allows employees to invest in company stock with numerous privileges including a discount, a look-back feature, and advantageous tax treatment. The simulation used empirical returns on company stock and other assets as a sample space from which it generated simulated returns to the ESPP and other assets. I used these simulated returns to estimate aggregate return distributions. Lastly, I compared these return distributions given employee preferences so I could posit optimal employee behavior.

I demonstrated through this simulation that, given reasonable assumptions, ESPP participation maximizes expected utility. Individuals who decline to participate at all effectively refuse compensation. However, individual employee risk-preferences and overall wealth should govern the holding period of the ESPP. This simulation suggests that long-term participation in an ESPP is often sub-optimal for employees because it severely undermines portfolio diversity. This is especially important because managerial concerns often lead employers to encourage long-term investment in company stock through ESPPs and other company stock ownership programs. In light of this study, perhaps employers should find alternative ways of aligning incentives.

As noted above, the methodology employed herein has several limitations. I did not perform sensitivity analysis on my tax rate assumptions. I did not test the implications of my data choices (i.e. the decision to exclude stocks with less than five years of returns or the decision to include all 40 years of available data instead of perhaps only 20 or 30 years of data). Finally, I did not perform this simulation in a lifecycle framework and thus

could not model the interaction between company stock returns and labor income. Future analyses could examine these unanswered questions.

Appendix 1: Tables and Figures

Table 1. Distribution of Total Gross Returns

	ESPP 6 Month	ESPP 18 Month	ESPP 36 Month	ESPP 60 Month	S&P 6 Month	S&P 18 Month	S&P 36 Month	S&P 60 Month	Portfolio 6 Month	Portfolio 18 Month	Portfolio 36 Month	Portfolio 60 Month
1%	1.1271	0.4611	0.3373	0.2911	0.8043	0.7340	0.6962	0.6835	0.9843	0.7539	0.7237	0.7383
5%	1.1271	0.6214	0.4956	0.4561	0.8749	0.8325	0.8210	0.8375	1.0078	0.8402	0.8181	0.8448
10%	1.1271	0.7323	0.6138	0.5862	0.9113	0.8874	0.8947	0.9332	1.0200	0.8955	0.8821	0.9201
25%	1.1273	0.9590	0.8874	0.9166	0.9703	0.9842	1.0316	1.1198	1.0398	1.0033	1.0190	1.0925
50%	1.1602	1.2911	1.3790	1.6032	1.0354	1.1007	1.2080	1.3744	1.0725	1.1528	1.2417	1.4062
75%	1.3612	1.8104	2.2927	3.1203	1.1019	1.2292	1.4155	1.6915	1.1616	1.3687	1.6154	2.0178
90%	1.6601	2.6145	3.9820	6.5159	1.1644	1.3576	1.6342	2.0435	1.2821	1.6796	2.2516	3.2682
95%	1.9084	3.3743	5.8474	10.9108	1.2035	1.4414	1.7822	2.2902	1.3779	1.9608	2.9231	4.8175
99%	2.5948	5.9224	13.7340	34.6889	1.2819	1.6150	2.1004	2.8406	1.6329	2.8682	5.6597	12.9391
Means	1.3058	1.5758	2.1873	3.8975	1.0369	1.1142	1.2428	1.4447	1.1215	1.2522	1.5228	2.1953
Std. Devs	0.3390	1.3367	5.3161	26.9187	0.1002	0.1867	0.2979	0.4566	0.1419	0.5027	1.8587	9.0985

Table 2. Distribution of Annualized Net Returns

	ESPP 6 Month	ESPP 18 Month	ESPP 36 Month	ESPP 60 Month	S&P 6 Month	S&P 18 Month	S&P 36 Month	S&P 60 Month	Portfolio 6 Month	Portfolio 18 Month	Portfolio 36 Month	Portfolio 60 Month
1%	0.2703	-0.4032	-0.3039	-0.2187	-0.3530	-0.1863	-0.1137	-0.073	-0.031	-0.172	-0.102	-0.059
5%	0.2703	-0.2718	-0.2086	-0.1453	-0.2346	-0.1150	-0.0636	-0.035	0.016	-0.110	-0.065	-0.033
10%	0.2703	-0.1876	-0.1502	-0.1013	-0.1696	-0.0765	-0.0364	-0.014	0.040	-0.071	-0.041	-0.017
25%	0.2707	-0.0275	-0.0390	-0.0173	-0.0586	-0.0106	0.0104	0.0229	0.0811	0.0022	0.0063	0.0178
50%	0.3461	0.1857	0.1131	0.0990	0.0720	0.0660	0.0650	0.0657	0.1503	0.0994	0.0748	0.0706
75%	0.8530	0.4854	0.3186	0.2556	0.2142	0.1475	0.1228	0.1108	0.3494	0.2328	0.1734	0.1507
90%	1.7559	0.8978	0.5850	0.4548	0.3558	0.2261	0.1779	0.1536	0.6438	0.4130	0.3107	0.2672
95%	2.6420	1.2497	0.8016	0.6128	0.4484	0.2760	0.2124	0.1802	0.8986	0.5666	0.4298	0.3695
99%	5.7330	2.2734	1.3948	1.0325	0.6433	0.3765	0.2807	0.2322	1.6663	1.0187	0.7821	0.6687
Means	0.7052	0.3542	0.2981	0.3127	0.0751	0.0747	0.0751	0.0764	0.2578	0.1618	0.1505	0.1703

Figure 1. Distribution of 6 month total gross returns

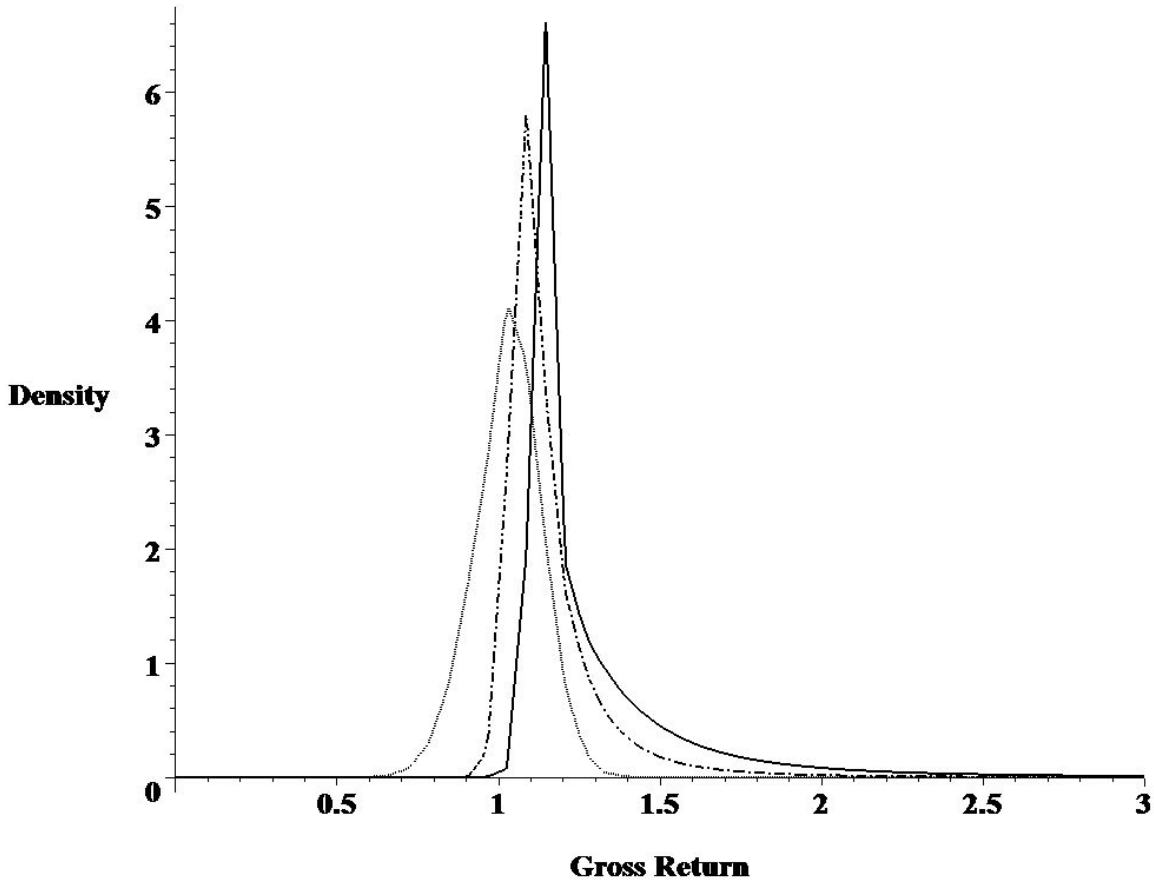

Legend
Solid = ESPP
Grey = S&P 500
Dashed = Portfolio

Figure 2. Distribution of 18 month total gross returns

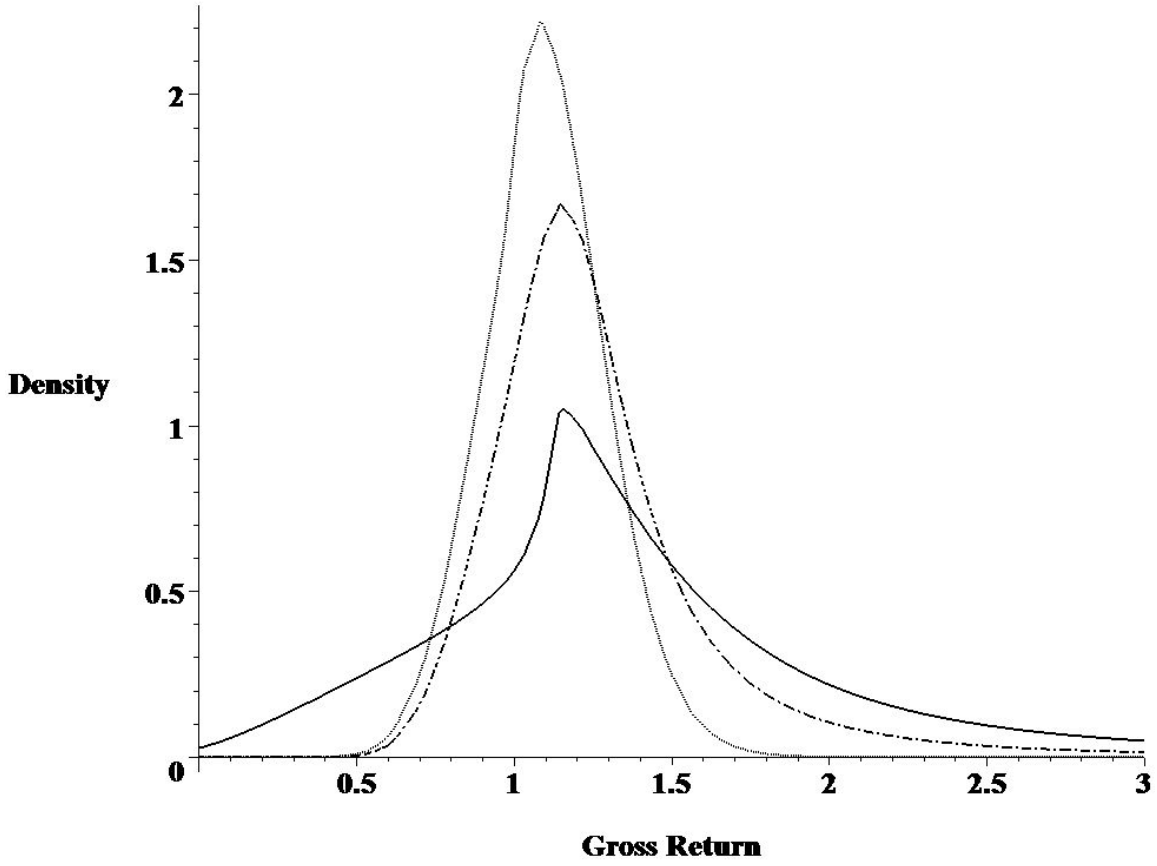

<u>Legend</u>
Solid = ESPP
Grey = S&P 500
Dashed = Portfolio

Figure 3. Distribution of 36 month total gross returns

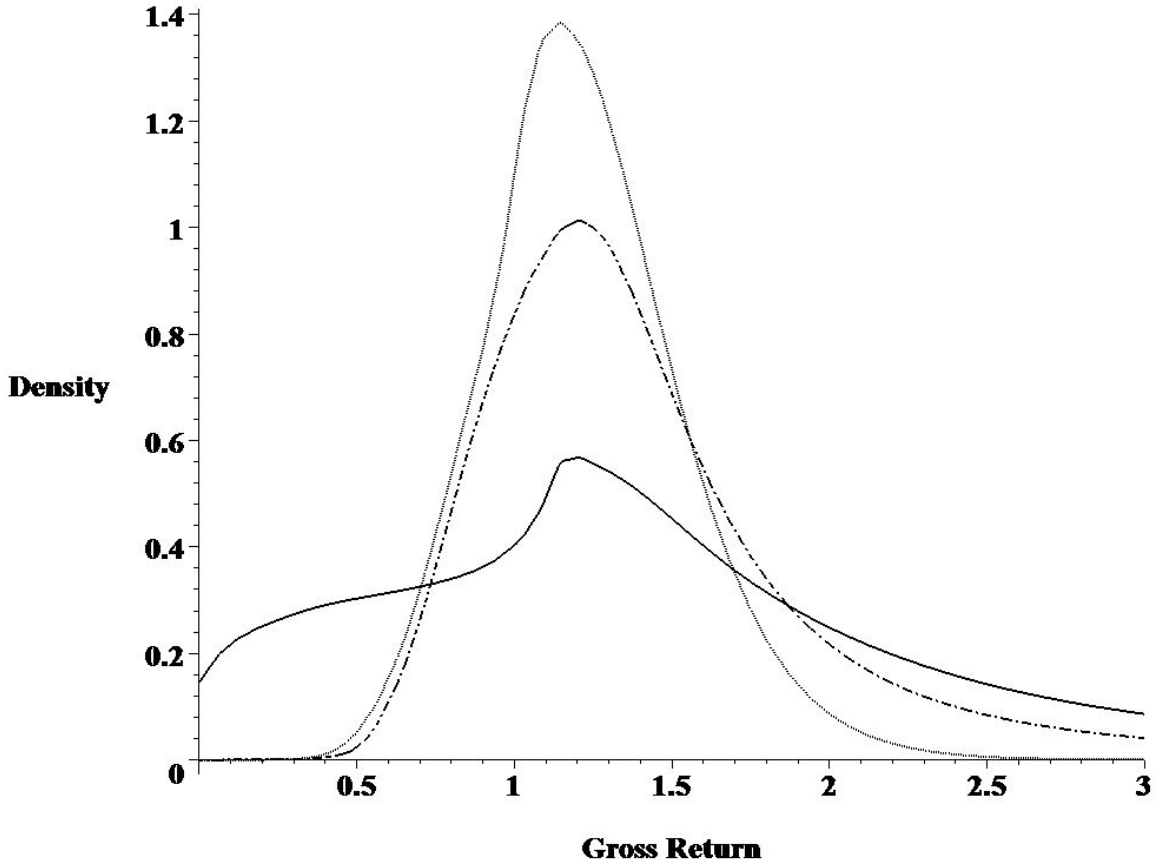

Legend
Solid = ESPP
Grey = S&P 500
Dashed = Portfolio

Figure 4. Distribution of 60 month total gross returns

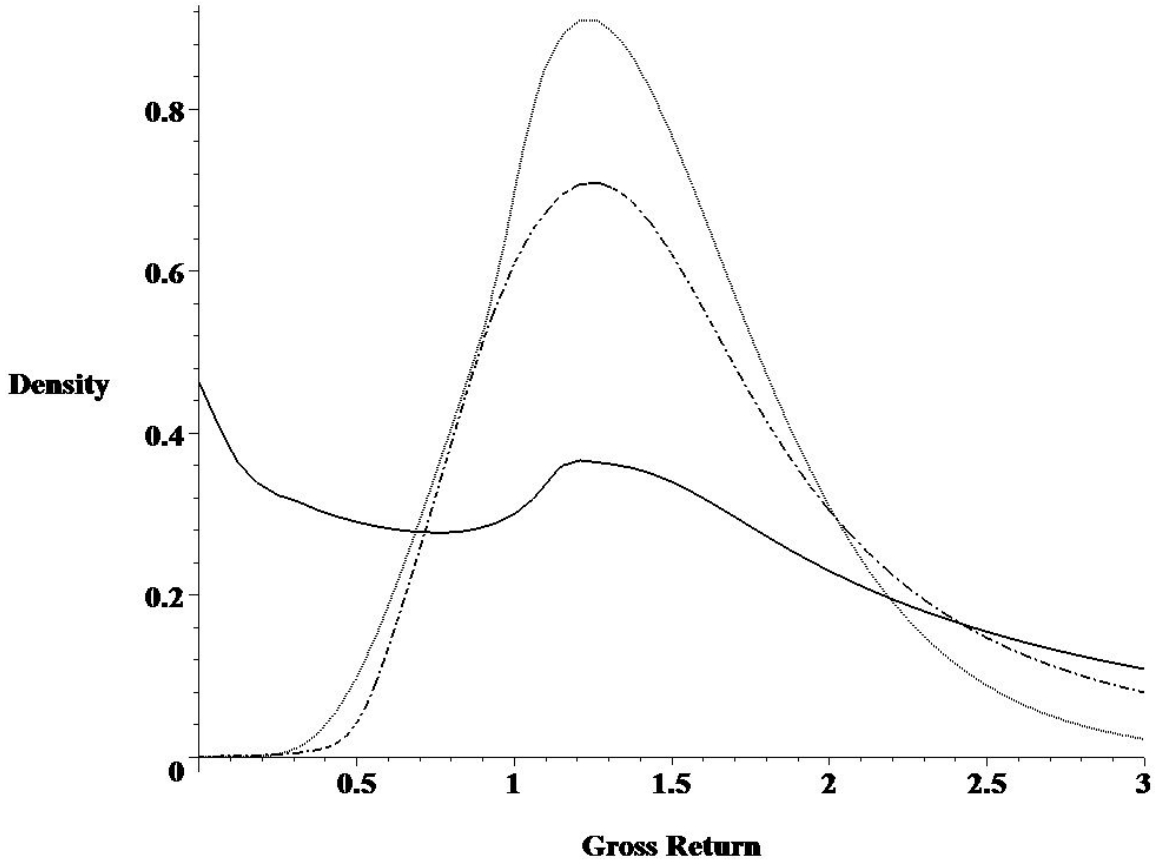

Legend
Solid = ESPP
Grey = S&P 500
Dashed = Portfolio

Table 3. Scenario 1 CRRA Analysis

CRRA Coefficients, Initial Investment = $3,000						
Strategy Comparison		**Other Wealth at 60 months**				
Risk-Neutral	**Risk-Averse**	**$0**	**$10,000**	**$50,000**	**$100,000**	**$300,000**
ESPP 0-6 -> S&P 7-60	**S&P 0-60**	>100	>100	>100	>100	>100
ESPP 0-18 -> S&P 19-60	**ESPP 0-6 -> S&P 7-60**	0.797	2.462	8.693	16.438	47.387
ESPP 0-36 -> S&P 37-60	**ESPP 0-18 -> S&P 19-60**	0.614	1.967	6.783	12.741	36.539
ESPP 0-60	**ESPP 0-36 -> S&P 37-60**	0.261	0.776	2.516	4.652	13.167

Table 4. Scenario 2 CRRA Analysis

CRRA Coefficients, Initial Investment = $3,000						
Strategy Comparison		**Other Wealth at 60 months**				
Risk-Neutral	**Risk-Averse**	**$0**	**$10,000**	**$50,000**	**$100,000**	**$300,000**
ESPP 0-6 -> Portfolio 7-60	**Portfolio 0-60**	>100	>100	>100	>100	>100
ESPP 0-18 -> Portfolio 19-60	**ESPP 0-6 -> Portfolio 7-60**	0.614	1.987	7.081	13.405	38.674
ESPP 0-36 -> Portfolio 37-60	**ESPP 0-18 -> Portfolio 19-60**	0.469	1.500	5.140	9.637	27.591
ESPP 0-60	**ESPP 0-36 -> Portfolio 37-60**	0.160	0.462	1.484	2.737	7.733

[i] Section 423 of the Internal Revenue Code

[ii] (1) The employee must hold the stock for at least one year after purchase date, and (2) two years after the beginning of the offering period

[iii] Brown, Jeffrey, Nellie Liang, and Scott Weisbenner, 2004, 401(k) Matching Contributions in Company Stock: Costs and Benefits for Firms and Workers, NBER Working Paper 10419, April