

Pay Versus Performance Is Here, Part 3: Valuing Stock Options for Compensation Actually Paid

October 26, 2022

On August 25, the SEC release the final [pay versus performance \(PvP\) rule](#). We've published [multiple articles](#) on the overall rule and core disclosure table. Now let's go narrow and deep on one particular dimension: valuing outstanding employee stock options for purposes of calculating compensation actually paid (CAP).

As you may know, the determination of CAP will require valuing outstanding options as of the beginning and end of covered fiscal years, as well as of the vesting dates. In order to do so, it'll be necessary to update the input assumptions, and consider whether to update the valuation methodology.

There's a lot of mystique around option valuation in this context. It's common knowledge that TSR-based awards with market conditions must be valued using Monte Carlo simulation and that vanilla RSUs are valued using the face value of the stock. But with options, it's not obvious what model to use and how to estimate the inputs to that model. Valuing a new 10-year at-the-money grant is quite different from valuing many grants that possess different moneyness levels and remaining lives.

We'll answer that question in this blog. Our stance is intentionally agnostic. It's better to focus on which methods fit different fact patterns than to default to simplistic one-size-fits-all answers. In the same spirit, we at Equity Methods deliberately keep our engagement pricing constant so that our clients can decide on the substantive merits.

Before we get into the methods, though, let's define the underlying problem.

The Core Problem: Exercise Behavior

Employee stock options (ESOs) differ from tradable stock options in that they're rarely held to maturity.¹ The crux of the ESO valuation problem is modeling early exercise behavior. In 2004 when FAS 123R (now ASC 718) was first released, there was an implicit preference for a lattice model. Shortly thereafter, the SEC released SAB 107 (now SAB Topic 14) and gave companies free reign. Today, over 80% of companies use the Black-Scholes model.

¹ Typical reasons for early exercise include liquidity needs, risk aversion and the inability to hedge risk as holders of tradable options can do, and departure from the company.

It's relatively straightforward to estimate the expected term for Black-Scholes. For companies with historical data, an analysis is performed on historical exercise behavior to infer how long, on average, options are held for. Companies without usable historical data typically use the SEC Simplified Method of SAB Topic 14 to formulaically calculate an expected term based on their vesting schedule and contractual term.

An expected term calculated at grant answers the question of how long an option is likely to be held assuming a broad distribution of stock prices. ASC 718-10-S99-1 explains that one reason historical data may not work is if there are “differences in terms of past equity-based share option grants” or “a lack of variety of price paths that the company may have experienced.”

For simplicity, suppose you're trying to value the following four grants that cliff vest after two years. On the date of grant, each option was assumed to have a 6-year expected term and a maturity of 10 years.²

| Grant | Remaining Term (years) | Moneyness |
|-------|------------------------|-----------|
| 1 | 9 | 60% |
| 2 | 8 | 120% |
| 3 | 7 | 130% |
| 4 | 6 | 80% |

Table 1: Four Grants that Cliff Vest After Two Years

To use its historical data to derive an expected term for each of these grants, the historical data would need to contain options that were at similar levels. For example, ideally the company could study a collection of options that had 6 years left on their life and an in-the-money ratio of 80%. After all, this is exactly what's underpinning a normal grant-date valuation insofar as you're studying the homogenous pool of prior options granted and their behavior after being granted with a 10-year life and starting at-the-money.

In reality, very few companies can sub-segment their historical data to find representative prior grants that had comparable remaining lives and comparable moneyness levels at those same points in time.

So this is the problem: How do we come up with a supportable expected term assumption for a wide-ranging portfolio of grants, all of which have different remaining terms and different moneyness levels?

Let's talk through a few approaches.

² The moneyness level is the ratio of the current stock price to the strike price. For example, a moneyness of 70% means that if the current stock price is \$7, the strike price is \$10.

Midpoint Method

This is a variation on the SEC Simplified Method of ASC 718-10-S99-1, though we can't call it the Simplified Method since that's prohibited for options that don't meet the plain-vanilla definition (including options not-at-the-money). Nonetheless, we can change the name and functionally adopt the same approach of calculating the midpoint between the weighted-average remaining time to vest and the expiration date. Basically, we split the difference between the earliest the option can be exercised and the latest it can be exercised.

Using our same options from above, let's now add the resulting expected term:

| Grant | Remaining Term (years) | Moneyness | Expected Term (years) |
|-------|------------------------|-----------|-----------------------|
| 1 | 9.0 | 60% | 5.0 |
| 2 | 8.0 | 120% | 4.0 |
| 3 | 7.0 | 130% | 3.5 |
| 4 | 6.0 | 80% | 3.0 |

Table 2: Four Grants that Cliff Vest After Two Years under the Midpoint Method

You can probably see the problem. The math works fine, but how in the world would grant #4 get exercised earlier than grant #3? It's got a lot more ground to cover to gain even a penny of value, whereas grant #3 is already in-the-money. As a spoiler alert, option exercise arguably has much more to do with moneyness than with raw time.

Contractual Term Method

Perhaps you're bewildered by the underwater options and figure you'll just use the remaining contractual term as a safe and conservative assumption. This is a second available method.

| Grant | Remaining Term (years) | Moneyness | Expected Term Approaches | |
|-------|------------------------|-----------|--------------------------|------------------------|
| | | | Midpoint Method (years) | Remaining Term (years) |
| 1 | 9.0 | 60% | 5.0 | 9.0 |
| 2 | 8.0 | 120% | 4.0 | 8.0 |
| 3 | 7.0 | 130% | 3.5 | 7.0 |
| 4 | 6.0 | 80% | 3.0 | 6.0 |

Table 3: Four Grants that Cliff Vest After Two Years under the Contractual Term Method

The problem is its reliability. For one, the remaining contractual term isn't a realistic assumption for grants 2 and 3 which are already handsomely in-the-money. More concerning, this approach directly clashes with the paradigm being asserted on the date of grant. How can

we say that a 10-year option that’s at-the-money will be exercised early, but a 7-year option that’s 30% in-the-money will be held until its expiration date?³

On page 64 of the final PvP rule, the SEC requires disclosure of the valuation assumptions used when those assumptions “materially differ from those disclosed at the time of grant.” Our concern with the contractual term method isn’t that we think it will trigger comment letters but that it will require dancing around the consistency problem noted above.

Elapsed Term Method

If the contractual term method concerns you because of its wobbliness on the underlying estimation paradigm, the elapsed term method can be a solution. Here we take the view that the expected term estimated at grant (which, in our example, is six years) is reliable and should be salvaged. So we simply subtract the amount of time that has passed.

| Grant | Remaining Term (years) | Moneyness | Expected Term Approaches | | |
|-------|------------------------|-----------|--------------------------|------------------------|----------------------|
| | | | Midpoint Method (years) | Remaining Term (years) | Elapsed Term (years) |
| 1 | 9.0 | 60% | 5.0 | 9.0 | 5.0 |
| 2 | 8.0 | 120% | 4.0 | 8.0 | 4.0 |
| 3 | 7.0 | 130% | 3.5 | 7.0 | 3.0 |
| 4 | 6.0 | 80% | 3.0 | 6.0 | 2.0 |

Table 4: Four Grants that Cliff Vest After Two Years under the Elapsed Term Method

But now we’re back to our original problem where we’re asserting that grant #4 is close to being exercised. It’s not. And grant #3 will certainly be exercised ahead of grant #4. There are other variations on this approach, but let’s jump to a different genre of tools: lattice models.

Lattice Model with Full Calibration

The customary way to implement a lattice model is to calibrate the full model on your historical data. If you have ample data, you’ll usually get a pretty robust result. If you don’t have much data, then you could calibrate the model using academic or other external data. In both cases, you still capture the primary benefit of adopting a model that produces expected term as an output.

³ Some companies use a contractual term for their options at grant for NEOs and other senior employees. For these, we believe that maintaining this makes the most sense.

The Basics of Lattice Models

Lattice models are predicated on the same underlying financial economic theory as Black-Scholes. A lattice model built using an expected term assumption will give the same value as Black-Scholes (unless there are dividends, where Black-Scholes is commonly known to have a limiting flaw). Lattice models, however, are more flexible. They can be built with more inputs whereas Black-Scholes unconditionally has six inputs—no more, no less.

A simple lattice model uses a target in-the-money ratio to explain exercise behavior, such as asserting that employees will exercise once the stock price reaches 200% of the strike price. Far more elaborate lattice models can also be developed.

The elegance of a lattice model is that it can be run the same on options without adjustment regardless of remaining term or moneyness. In effect, it produces the expected term as an output instead of requiring it as an input. It models exercise behavior in the context of relevant parameters and outputs an expected term. In other words, it sidesteps the need to come up with an expected term as we try to do in the tables above.

Assuming a 50% volatility, 0% dividend yield, and 4% interest rate, our table continues to expand.

| Grant | Remaining Term (years) | Moneyness | Expected Term Approaches | | | |
|-------|------------------------|-----------|--------------------------|------------------------|----------------------|--------------------------|
| | | | Midpoint Method (years) | Remaining Term (years) | Elapsed Term (years) | Full Calibration (years) |
| 1 | 9.0 | 60% | 5.0 | 9.0 | 5.0 | 7.30 |
| 2 | 8.0 | 120% | 4.0 | 8.0 | 4.0 | 3.35 |
| 3 | 7.0 | 130% | 3.5 | 7.0 | 3.0 | 2.48 |
| 4 | 6.0 | 80% | 3.0 | 6.0 | 2.0 | 4.47 |

Table 5: Four Grants that Cliff Vest After Two Years under a Lattice Model with Full Calibration

The numbers look intuitive, but where did they come from? We used a simple Hull-White model where exercise is a function of reaching a particular in-the-money barrier, in this case 170% of the strike price. To get to a term, we then determined what term we could plug into a Black-Scholes in order to compare the model. Well, where did that come from? It came from an analysis of the company's historical data.

Of course, we expect that companies using a lattice model for their grant date fair values will continue to do so for their PVP calculations.

All this is fine, and ASC 718 is explicitly clear in ASC 718-10-55-17 and 55-20 that it's acceptable to use different models for different valuation circumstances. But this nonetheless requires added disclosure in the PVP footnotes since it represents a deviation in the assumptions from those used at grant.

Lattice Model with Calibration to Grant Date Fair Value

The last method, and the one we're seeing adopted the most, involves a convenient technique offered in Footnote 55 of the original FAS 123R, now ASC 718-10-55-30. This method uses the Black-Scholes grant-date fair value to back-solve for the appropriate lattice model input (the in-the-money ratio). This way, the lattice model's assumptions about exercise are directly, mechanically linked to those used in Black-Scholes at the point of grant. Importantly, we won't see a value shift due solely to the model change. We're simply using the lattice model to more objectively compute the expected term for later valuations in a way that's consistent with the grant-date assumptions.

To do the math, we need one more input, the grant-date fair value in Black-Scholes. For simplicity, let's assume that each of our two-year cliff vesting options was granted at \$100 and had the same other inputs to yield a value of \$52.42. This yields an in-the-money barrier of 2.27 and lets us fill out our final column in the table:

| Grant | Remaining Term (years) | Moneyness | Expected Term Approaches | | | | |
|-------|------------------------|-----------|--------------------------|------------------------|----------------------|--------------------------|-----------------------------|
| | | | Midpoint Method (years) | Remaining Term (years) | Elapsed Term (years) | Full Calibration (years) | Reverse Footnote 55 (years) |
| 1 | 9.0 | 60% | 5.0 | 9.0 | 5.0 | 7.30 | 7.96 |
| 2 | 8.0 | 120% | 4.0 | 8.0 | 4.0 | 3.35 | 5.12 |
| 3 | 7.0 | 130% | 3.5 | 7.0 | 3.0 | 2.48 | 4.25 |
| 4 | 6.0 | 80% | 3.0 | 6.0 | 2.0 | 4.47 | 5.16 |

Table 6: Four Grants that Cliff Vest After Two Years under a Lattice Model with Calibration to Grant Date Fair Value

We used the same volatility and interest rate in every valuation, but as a practical matter, these would likely vary simply due to the passage of time. How would this method get disclosed? We envision language such as:

We used a binomial lattice model to value outstanding options included in the measurement of compensation actually paid. This allows us to estimate the value of each option in the context of each option's remaining contractual life and in-the-money ratio at the point of measurement. To maintain consistency with our approach toward grant-date fair value estimation under ASC 718, we inferred the binomial lattice model's exercise barrier from the Black-Scholes grant-date fair value. That input is an assumed in-the-money ratio at which exercise will occur, which is comparable to the model used throughout the illustrative examples in ASC 718.

Of course, there are other approaches to using a lattice model and estimating the expected term assumption. These include using data from academic research, actuarial research, and other tools for calibrating either a lattice model or Black-Scholes. But rather than building an excessively exhaustive list, let's switch gears and cover how to compare and contrast the different approaches on the table.

How to Choose an Appropriate Technique

Our view is that you should begin with your fact pattern and bring all the available tools to the table. We've chosen not to vary our fees by approach, because technology lets us capture similar efficiencies and we think the decision should fit each company's own context.

More often than not, the most effective methods turn out to be the Elapsed Term Black-Scholes and Reverse Footnote 55 Lattice. The Black-Scholes approach is crude but incredibly simple. The lattice approach is more complex but also much more precise and robust. We generally don't prefer in-between approaches that add complexity for its own sake but still suffer from the same inherent flaws as the simplest models.

Let's look at two factors that can aid decision-making. First, do you pay a dividend and are your options in-the-money? If so, this may trigger one of Black-Scholes' known flaws. Second, is there a wide spread of in-the-money levels? The more scattered they are, the more you'll benefit from a more reliable technique. We'll cover these topics next.

Lattice vs. Black-Scholes: Dividends and Moneyness

A commonly understood problem with Black-Scholes is the manner in which it treats dividends. This problem is accentuated when an option is in-the-money.

Let's start with an example considering an option at its vesting date:

- **Current stock price: \$40**
- Strike price: \$20
- Expected term (years): 6

- **Volatility: 30%**
- Risk-free interest rate: 3%
- Dividend yield: 4%

This option's intrinsic value is \$20. Using these assumptions, Black-Scholes produces a fair value of \$16.50 because it assumes the option will be held to maturity, and the option holder loses the value of six coming years of dividends. Indeed, the holder could exercise the option today

and get \$20, which sets a floor for the value and means that Black-Scholes would not be ASC 718-compliant in this case.⁴

This is why we encourage you to be leery of simplistic, one-size-fits all approaches. In this case, a lattice model provides a result that's less likely to trigger a comment letter or other form of scrutiny. It can be a simple lattice model, as long as it doesn't trigger one of the known limitations of Black-Scholes.

Lattice vs. Black-Scholes: Year-over-Year Swings

The impact of model selection on valuation is fairly straightforward:

- If the stock price goes up, we expect the value of the option to go up, and conversely for a falling stock price
- Further, options are leveraged instruments, meaning that if the stock price goes up (say by 10%), we expect the option value to move by more than this
- On the other hand, as time passes, the option value goes down. The remaining chance for the stock price to move up (or alternatively, the insurance value of being able to walk away in a down market) reduces. In option parlance, this is called theta bleed, theta being the term for the movement in option value over time⁵

As discussed above, one would intuitively expect an increasing price to result in a shorter holding period. As a result, this shorter term mitigates the change in value from the stock price alone. Lattice models allow us to capture this effect and reflect smaller year-over-year value variance.

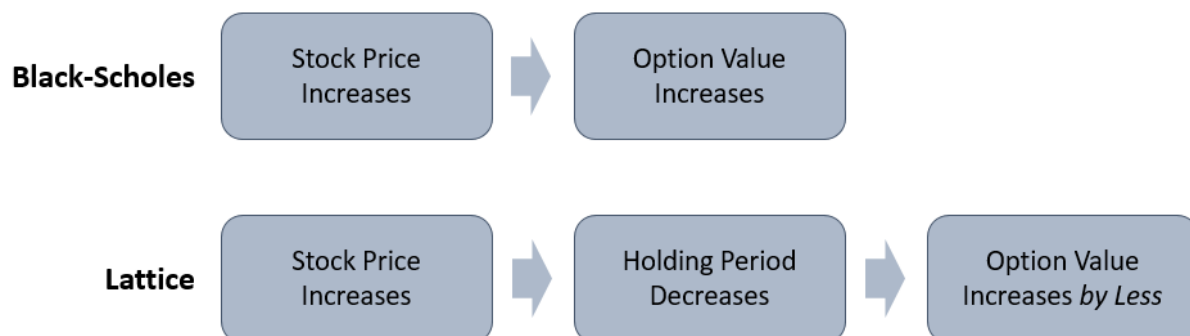


Figure 1: Lattice vs. Black-Scholes Models

⁴ The Cox, Ross, Rubenstein (CRR) lattice model takes into account optimal exercise incorporating future dividends, and can be implemented into the aforementioned lattice with a stock price hurdle to capture optimal as well as employee “sub-optimal” early exercise.

⁵ We note that theta bleed would also be a great name for a punk rock band and/or action movie franchise.

For a stock price decrease, the impact is reversed. In order to show the impact, we show the results comparing two models.⁶

- The first is a Black-Scholes model using the elapsed term method. This assumed a seven-year expected term at the grant date, and at the end of a year, it assumes six years remaining.
- The second is a lattice model which assumes exercise at the end of the period, or if the stock price rises to 150% of the strike. Because higher stock prices are closer to the hurdle, this will shorten the term in these cases.

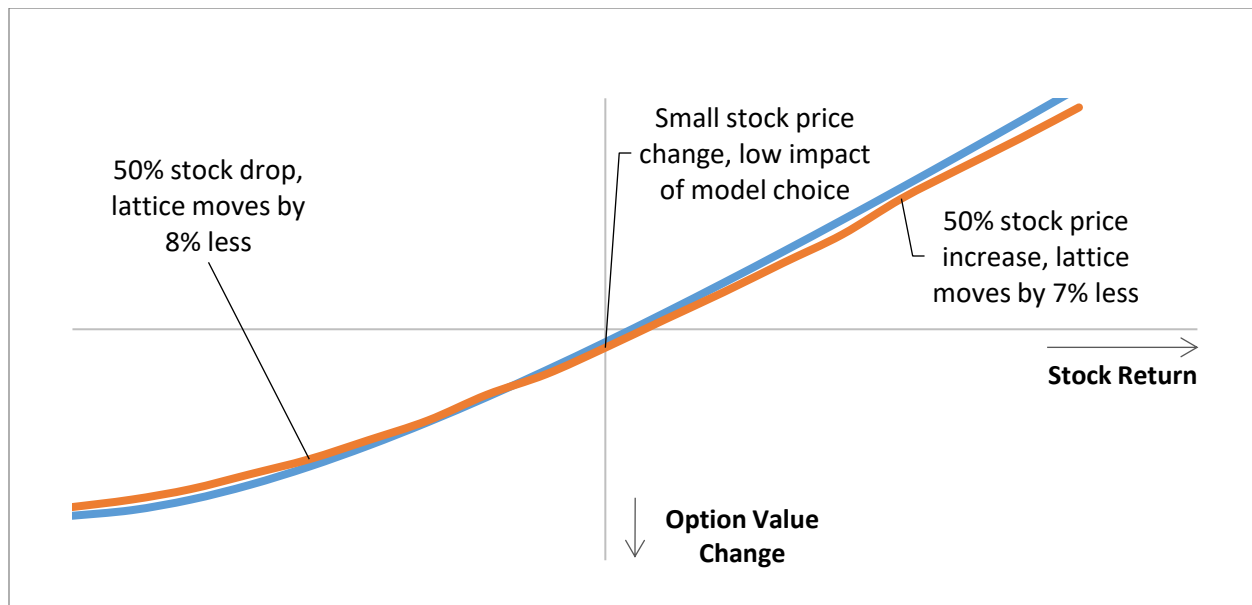


Figure 2: Option Value Moves under Lattice vs. Black-Scholes Models

The results show exactly what we anticipated. If the stock price stays flat, we see a low drag on the option value from the passage of time, approximately 3%. The stock needs to go up around 5% to keep the value neutral.

On both sides of the curve, we see the lattice value results in a smaller impact than the Black-Scholes value. In our example, we see this is frequently in the 7-8% range, meaning this could show a definite impact on the disclosed figures, especially for companies with a large portion of their compensation as options.

Of course, your mileage may vary. But the point is that in some contexts lattice models may give a strategically more desirable result (while also being more accurate). The cons could outweigh that benefit, which is why it's good to consider the full spectrum of alternatives.

⁶ We assume the option is at the money, stock and strike is at \$0 at the grant date, and the term is a full year. Our option cliff vests in 4 years and expires in 10. Volatility is 35%. Risk-free rate is 2%. Our stock does not pay dividends. Post-vest termination rate in our lattice model is 2%.

Final Thoughts

Equity award valuation is proving to be the most complex piece in the calculation of CAP. TSR awards require separate Monte Carlo models that are built to capture the distinct terms of each award. Options are different because they're all the same (e.g., 10-year term) and are simply at different points in their life and at different moneyness levels, referred to as being "on the run."⁷

Given this homogeneity in option terms, the best valuation method is the one that can mass-process the entire portfolio of options subject to the PvP calculation, yield reliable results for each option, and require minimal to no hunting and pecking through individual results because they're contradictory or unreliable. The model that accomplishes that could be Black-Scholes. On the other hand, Black-Scholes could also be the source of unsupportable results, such as some of the cases we reviewed in this paper.

That's why our approach with clients is to look at the data first, then evaluate the right tradeoff between model rigor and model simplicity. We make it our clients' decision, we offer our recommendation and perspective, and we intentionally don't vary our fees by the decision. The goal is for you to develop a methodology that's efficient, scalable to all your options subject to the CAP calculation, and supportable, with a clean disclosure.

If you have any questions about this new rule and how it impacts your company, feel free to reach out.

Takis Makridis

President & CEO

takis.makridis@equitymethods.com

480-428-1203

David Outlaw

Managing Director

david.outlaw@equitymethods.com

480.428.3305

Josh Schaeffer, PhD

Managing Director

josh.schaeffer@equitymethods.com

480.428.1175

⁷ While we're not aware of a band named after Theta Bleed, it's conceivable that REO Speedwagon and/or Paul McCartney and Wings had some exposure to valuing options on the run.

About This Document

The content of this document is for general information purposes only. It is distributed under the understanding that it does not constitute professional advice or services, and therefore it should not be used or relied upon as the basis for any action or decision that may affect you or your business. Equity Methods accepts no liability related to the content of this publication.

Additionally, this document is subject to our general [website terms of use](#). Any effort to infringe on the wording or overall content is a violation of our intellectual property and will be prosecuted.

www.equitymethods.com

Equity Methods • 17800 N. Perimeter Drive, Suite 200 • Scottsdale, AZ 85255 • 480.428.1200 • info@equitymethods.com