

Agency Under Leverage

Why the Best Capital Structures Protect the Ability to Move

By Robert E. Tucker

Grey Matter Partners

January, 2026

Executive Summary

Organizations already allocate capital using counterfactual logic. Downside counterfactual spending such as insurance, compliance, cybersecurity, and risk management is explicit, normalized, and protected because failure is visible, costly, and often existential, even though its value lies in outcomes that never occur. This logic is deeply embedded in how boards, lenders, and sponsors think about risk and capital discipline.

The same logic applies on the upside, but it is treated far less rigorously inside operating companies. Markets explicitly price upside counterfactuals through liquidity, equity, and options that preserve future opportunity, while firms handle comparable investments in leadership capacity, governance, M&A awareness, and decision bandwidth informally and defensively. This paper introduces **Possibility Creation Value (PCV)** as a framework for naming and pricing these investments. PCV does not create success; it preserves choice. In leveraged environments, its absence often becomes visible only when a refinancing, covenant breach, or market shock collapses the outcome space into a forced decision. A private equity case study shows that explicitly budgeting for PCV alongside downside protection materially widens strategic options and improves outcomes when markets tighten.

Grey Matter Partners
rob@greymatterpartners.com

Framing the Problem

Most organizations like to think of themselves as rational actors. They believe they allocate capital based on evidence, discipline, and experience. Yet many of the most important decisions leaders make are not grounded in what has already happened, but in what plausibly could have happened, or still might. This is especially true under uncertainty, where the cost of being wrong is asymmetric and the future does not repeat cleanly. The analysis that follows begins with this observation: some of the most consequential business decisions are justified not by realized outcomes, but by alternatives that never materialize.

1. The accepted value of counterfactual reasoning

Most serious business decisions are not made by looking only at what happened. They are made by asking what *might have* happened, and what would have happened *instead* if a different choice had been made. That kind of thinking is called counterfactual reasoning, but the idea itself is simple. You imagine another version of the future and decide whether it is better or worse than the one you are trying to create.

Companies do this all the time, even if they do not use the term. When a firm buys insurance, it is not buying an outcome. It is buying protection against a future that never arrives. When a company invests in cybersecurity, most days nothing happens at all. That is the point. The value comes from the breach that did not occur, the lawsuit that was never filed, the factory that did not burn down.

This kind of spending feels normal because the logic is familiar. A small, predictable cost is exchanged for protection against a large, uncertain loss. No one expects proof that the bad thing would have happened. The possibility is enough. The reasoning is widely accepted, deeply embedded, and rarely questioned.

What matters is not that the outcome is invisible, but that the alternative outcome was plausible and expensive. Counterfactual reasoning allows organizations to act responsibly in the face of uncertainty, even when success looks like nothing happening at all.

2. Downside counterfactual spending (DCS) is explicit and normalized

Because this way of thinking is so familiar, organizations have built entire budgets around it. Downside counterfactual spending is not hidden or controversial. It is written into annual plans, reviewed by boards, and defended by auditors. No one asks whether it is real. The only question is whether it is sufficient.

Insurance is the clearest example. Companies insure buildings that never catch fire, executives who are never sued, and shipments that arrive safely every day. Compliance

works the same way. Firms spend heavily to meet regulations they never violate. Cybersecurity teams work around the clock to stop attacks that, if they are successful, become news precisely because they are rare.

Risk management extends this logic across the organization. Redundant systems, backup suppliers, legal reviews, internal controls — all are designed to prevent problems that most employees will never see. These expenses are not treated as optional or experimental. They are considered the cost of being serious.

What makes downside counterfactual spending different is not the certainty of the threat, but the clarity of the failure. Everyone understands what a lawsuit looks like. Everyone understands a data breach or a factory shutdown. Because the bad outcomes are easy to imagine and expensive to fix, spending to prevent them feels responsible, even when the money produces no visible result.

Over time, this has made downside counterfactual spending routine. It is planned, justified, and protected. No one calls it speculative. It is simply how grown-up organizations behave.

3. Upside counterfactuals instruments exist, but are treated informally inside organizations

The strange thing is that this logic does not stop at preventing harm. It already exists on the other side as well. Financial markets, in particular, are very comfortable paying for futures that might never arrive. Investors hold stock not because something good has already happened, but because it might. They buy long-dated options knowing most will expire worthless, because a few preserve access to outcomes that would otherwise be unreachable. They hold cash even when it earns little, because having money available keeps choices open.

Outside the walls of a company, this kind of spending is treated as disciplined and rational. It is priced, modeled, and debated in public. Inside organizations, however, the same logic is rarely acknowledged. Budgets for making good outcomes possible are not usually evaluated with the same care as budgets for preventing bad ones. Instead, they are discussed in softer terms: judgment, intuition, flexibility, or “keeping our options open.”

This creates an odd mismatch. The tools to reason about upside counterfactuals clearly exist, and markets use them every day. But inside firms, decisions about preserving opportunity are often made without a shared framework or a common language. They are justified one meeting at a time, vulnerable to mood, personality, and short-term pressure. Good things that might have happened are rarely examined with the same rigor as bad things that did not.

As a result, organizations are comfortable paying to stop failure, but uneasy paying to make success possible. The difference is not logic. It is how clearly the logic has been named.

Liquidity, equity ownership, and long-dated call options show that markets explicitly price upside counterfactuals, while corporate budgets rely on informal, intuitive, and emotionally driven decisions for the same logic.

4. A rigorous framework for upside counterfactual value in organizations

To balance the way organizations think about preventing harm, they need a comparable way to think about enabling success. German economic writing uses the word *Ermöglichungswert*, which translates most directly as “the value of making possible.” It captures something English usually talks around but rarely names. Possibility-Creation Value, or PCV, is a way to give that idea a clear place in decision making.

PCV refers to investments that do not guarantee good outcomes, but make good outcomes feasible. Liquidity is the simplest example. So are leadership structures that can actually decide, governance that allows speed instead of paralysis, senior hires who can carry weight, exploratory investments that open paths rather than close them, awareness of acquisition targets before a deal window opens, reputations that buy trust, and the protection of decision capacity under pressure. None of these produce results on their own. They create the conditions in which results can happen.

This is why PCV can look deceptive. Many of its instruments appear inactive. Holding cash looks like waiting. Owning equity looks passive. Buying options looks like paying for nothing. Inside organizations, maintaining slack, protecting leadership attention, or investing ahead of clear need can look indulgent. But that appearance misunderstands their purpose. These investments are not meant to act. They are meant to preserve agency so that action is possible when the moment arrives.

Downside counterfactual spending protects survival. PCV protects freedom of movement. Both operate under uncertainty. Both rely on outcomes that may never occur. Treating one as rigorous and the other as optional leaves organizations lopsided. PCV does not replace discipline around risk. It completes it.

Drawing on *Ermöglichungswert* (“the value of making possible”), Possibility-Creation Value (PCV) names the upside counterfactual value created by investments that preserve or expand future feasibility, including liquidity, leadership structure, governance, senior hiring, exploratory investment, M&A awareness, reputation, and protected decision capacity. PCV instruments such as liquidity, equity, and options are best understood as preserving future agency rather than producing immediate action, and their apparent inactivity reflects their purpose rather than their failure.

5. **PCV should be discounted to reflect real asymmetry (dealing with objections)**

A serious framework for Possibility-Creation Value has to accept an uncomfortable fact. Not all risks are equal. In many businesses, the downside is brutal and final, while the upside is gradual and competitive. A refinery explosion, a regulatory shutdown, or a liquidity crisis can end a company outright. No single upside opportunity carries that kind of weight. Ignoring that asymmetry would be irresponsible.

But acknowledging asymmetry does not mean abandoning PCV. It means pricing it correctly.

Rational decision making already does this elsewhere. Insurance premiums rise when the stakes are higher. Capital requirements increase when failure would be catastrophic. The same logic applies to possibility creation. When upside opportunities are modest, uncertain, or easily copied, investments that preserve optionality should be discounted. They should clear a higher bar. They should be smaller, more focused, and more tightly defended. What they should not be is waved away as unserious.

The mistake many organizations make is confusing discounting with dismissal. They treat the fact that upside is incremental as proof that investments in future agency are indulgent. That is a category error. Incremental does not mean unimportant. Over time, small advantages compound. Slightly better timing, slightly better judgment, and slightly more freedom to act can separate firms that endure from those that merely survive.

A rational PCV framework would therefore ask different questions, not softer ones. How large is the feasible outcome space being preserved. How likely is it that this space will matter under stress. How reversible is the investment if conditions change. How much agency is actually being protected. These are not philosophical questions. They are the same questions risk managers already ask about downside exposure.

In practice, this means PCV is not a blank check. It is a measured allocation. When downside risks are existential and upside opportunities incremental, PCV should be weighted by stricter multipliers, scaled to context, and revisited regularly. Treated this way, it becomes neither reckless nor vague. It becomes disciplined.

Organizations do not fail because they miscalculate risk alone. They also fail because, at the moment when a decision mattered, they had no room to move. PCV is about making sure that moment still exists.

6. **A balanced case study: DCS and PCV in practice**

Consider a private equity owned company acquired in 2021 near the top of the market. The asset was purchased for \$1.6 billion, financed with aggressive leverage priced in a low-rate environment that no longer exists. Revenue is solid at \$420 million annually,

EBITDA is \$120 million, but interest expense has climbed to \$95 million. The company is not distressed, but it is not comfortable. There is no margin for error.

The owners understand downside risk very clearly. They approve \$22 million a year in downside counterfactual spending. That includes insurance across property, D&O, and cyber at \$6 million. Compliance and regulatory monitoring cost another \$5 million. Cybersecurity and systems redundancy add \$4 million. Legal reserves, audit depth, and contingency planning account for the remaining \$7 million. None of this produces growth. All of it protects survival. The spend is explicit, reviewed quarterly, and defended without embarrassment.

Now look at the other side.

The company holds \$45 million in liquidity. That cash earns little, but it preserves optionality. It allows the firm to fund a tuck-in acquisition quickly, to bridge a refinancing delay, or to withstand a covenant breach without handing control to lenders. The opportunity cost of that liquidity is roughly \$4 million a year. No one celebrates it, but no one questions it either.

Beyond liquidity, however, PCV spending becomes fuzzy.

The firm has one overstretched operating partner covering four portfolio companies. There is no dedicated senior executive focused on refinancing strategy, lender negotiations, or scenario modeling. M&A monitoring is ad hoc, handled when time permits. Board agendas are crowded with reporting, leaving little protected time for forward-looking decisions. Reputation management and lender relationships depend almost entirely on the CEO, who is already managing day-to-day operations under stress.

In practice, the company spends perhaps \$3 million a year on these possibility-preserving activities, mostly buried in overhead. That figure is never discussed as a category. It is the first place cuts are suggested when pressure rises.

A disciplined PCV framework would rebalance this without pretending upside is limitless.

The firm could rationally allocate an additional \$6 to \$8 million annually to PCV. That money would fund a senior capital markets executive, not a rainmaker, but someone whose job is to preserve refinancing options. It would support systematic M&A awareness so acquisition opportunities are known before they are urgent. It would protect board time explicitly reserved for forward strategy rather than backward reporting. It would strengthen lender and rating-agency communication before negotiations turn adversarial.

None of this guarantees success. What it does is widen the feasible outcome set. Instead of two paths, refinance or default, the firm now has several. Amend and extend. Partial

asset sale. Minority recapitalization. Opportunistic acquisition that changes scale. These options exist only if the groundwork is laid in advance.

Importantly, the spending remains asymmetric. The company still spends more on downside protection than on possibility creation. That reflects reality. But the PCV spend is no longer accidental or apologetic. It is deliberate, priced, and defended using the same logic as insurance.

For a leveraged asset in a tight market, this balance matters. Most failures in this situation do not come from missing data or bad intentions. They come from having no room to maneuver when conditions shift. Downside counterfactual spending keeps the lights on. Possibility-Creation Value makes sure there is still a door to walk through when the moment arrives.

7. Call to action

The point of this framework is not to encourage optimism or loosen discipline. It is to bring the same seriousness to upside counterfactuals that organizations already bring to downside risk. On the income statement, downside counterfactual spending is easy to find. It lives in insurance premiums, audit fees, cybersecurity budgets, and compliance costs. These line items are protected because everyone understands what happens when they are cut.

Possibility-Creation Value is different. It rarely appears as a named category. It is scattered across salaries, retained liquidity, advisory relationships, board structure, and time. Because it is fragmented, it is easy to underfund and easier still to defend only after the fact. When pressure mounts, PCV is often the first thing reduced, precisely because its value is not immediately visible.

That is a mistake rooted in accounting, not economics.

The same expected-value logic that justifies spending millions to prevent rare failures also justifies spending measured amounts to preserve future agency. On a P&L, both are costs. In reality, both are investments in keeping the company alive long enough to make choices. One protects against collapse. The other protects against constraint.

The discipline required is straightforward. Organizations should explicitly identify PCV, assign it a budget, and subject it to the same scrutiny as downside risk. Not to inflate it, but to defend it. To ask whether it is sized correctly, whether it is focused on the right constraints, and whether it meaningfully expands the feasible outcome space.

Firms that do this will not eliminate risk. They will eliminate surprise. They will enter difficult periods with more than one move available. In uncertain environments, that difference is often decisive.

The recommendation is simple. Treat Possibility-Creation Value as real. Price it deliberately. Protect it when pressure rises. Companies already know how to think this way. They just need to apply it symmetrically.