

FOUNDATIONS PHASE

Beyond MIP vs MEP: the terminology, the sizing, and the conversations no one has

Author Aaron Roes Reading time 13 min read Source mipdesk.com/insights-mip-vs-mep.html

Walk into any private equity deal and you will hear at least three terms used to describe the same thing: MIP, MEP, sweet equity. The labels matter less than people think. The design choices behind them matter far more than people discuss.

This article does two things. First, it walks through the terminology that surrounds a management incentive programme, what each term means, how they relate, and where the differences are real versus cosmetic. Second, it tackles the question that actually determines whether a programme works: how to size the pool, and how to design a pool that survives the holding period.

There are no benchmark percentages in what follows. There is a reason for that: published benchmarks for MIP pool size are scarce, and the numbers that circulate informally tend to reflect one fund's recent deals rather than any market truth. The honest answer is that the right pool depends on the situation. What this article gives you is a structured way to think about it.

The terminology jungle: what the labels actually mean

In a typical PE deal you will hear half a dozen terms thrown around, often interchangeably. They are not random, each refers to a specific piece of the structure, but the labels overlap in ways that confuse new readers and, occasionally, experienced ones too. Here is the lay of the land, walked through in the order that the concepts enter a deal.

CAPITAL INSTRUMENTS

When a PE fund manager acquires a portfolio company, the capital structure typically includes three layers: senior debt at the top, preferred shares in the middle, and ordinary shares at the bottom. The PE fund manager invests a combination of preferred shares and a small slice of ordinary shares as an integrated package, this is the institutional strip, or simply the strip. Management invests in ordinary shares only, pure equity with no preferred component attached. These pure ordinary shares are called sweet equity precisely because they sit at the bottom of the stack and benefit from the leverage created by the preferred shares sitting above them.

THE PROGRAMME LABELS

The contractual framework through which management acquires and holds these ordinary shares goes by several names:

TERM	WHAT IT EMPHASISES	WHERE YOU HEAR IT
MIPManagement Incentive Plan	The incentive mechanism, motivation, alignment, performance	UK and broader European PE market. The most widely used term.
MEPManagement Equity Plan	The equity holding itself, ownership, share class, legal structure	Used interchangeably with MIP. Slight preference among practitioners focused on the structural side.
Sweet equity	The economic instrument, pure ordinary shares with maximum gearing	Informal shorthand for what management actually holds. Used more by deal teams than HR or legal.
Management participation	Continental European phrasing for the same concept	DACH and Benelux corporate counsel.

The honest answer to "MIP or MEP, what is the difference?" is: not much, in most contexts. They describe the same underlying mechanism. Where the labels do carry weight is in what they signal about the programme's design philosophy. A fund manager who consistently uses "MIP" tends to frame the programme around behaviour and alignment, the motivation to outperform. A fund manager who consistently uses "MEP" tends to frame the programme around ownership and governance, what management holds and under what rules. The same programme can be designed either way; the framing influences how the documentation is drafted and how the conversation with management is conducted.

THE MECHANISMS ATTACHED TO A PROGRAMME

Beyond the capital instruments and the programme label, several mechanisms shape how the programme actually works. Each is worth a brief mention here and will be unpacked in later articles in this series.

- Vesting, the schedule that determines how much of your ordinary share allocation belongs to you economically over time. Most programmes vest progressively over the holding period, often with a cliff in the early quarters during which no vesting accrues.
- Leaver provisions, what happens to your shares if you leave before exit. Three categories typically apply: good leaver, bad leaver, and early leaver. The classification determines the repurchase price.
- Ratchets, a mechanism that increases management's share of the proceeds at exit if certain performance thresholds are exceeded. Not standard in every deal; common in some sectors and structures.
- Hurdles, a return threshold the PE fund manager must achieve before management's economic interest fully activates. In strip-and-MIP structures, the hurdle is typically implicit in the preferred share return rather than a separate mechanism.
- Drag-along and tag-along, provisions that ensure all shareholders move together at exit, preventing minority holdouts from blocking a sale.

Each of these is a design choice. They are not given. The combination of choices is what makes one programme work and another disappoint everyone involved.

How the pieces fit together: a programme over time

Below is the lifecycle of a typical MIP, showing when each concept enters the picture. Read it left to right.

Sizing the pool: the question everyone asks, and nobody answers cleanly

Once the terminology is settled, the practical question that follows is: how big should the pool be? What percentage of the ordinary shares should be reserved for management?

Ask five practitioners and you will get five answers, usually expressed as a range, qualified with "it depends," and not always backed by data. Published benchmarks exist for some sub-segments of the market, but they vary widely by deal size, sector, geography, and the leadership intensity of the business. A number that works for a €50m healthcare deal will not necessarily work for a €500m industrial deal.

Rather than offering a benchmark, what follows is a framework for thinking about the question. The variables that should drive pool size, the spectrum of trade-offs, and an honest opinion about where the market tends to land.

THE VARIABLES THAT DRIVE SIZING

A handful of factors should shape the answer in any given deal:

- Deal size and equity cheque. A €1m allocation on a €50m deal feels generous; the same on a €500m deal feels symbolic. Pool sizing is a percentage decision, but the absolute amounts matter for whether the programme is materially motivating to individuals.
- Number of participants. A pool spread across five people produces a different individual outcome than the same pool spread across thirty. The wider the participation, the larger the pool needs to be to deliver meaningful gearing to the most senior participants.
- Leverage in the capital structure. A more leveraged structure (higher preferred relative to ordinary equity) generates more gearing on the sweet equity, meaning a smaller pool can deliver larger absolute returns at exit. Less leveraged structures produce flatter outcomes and may justify a larger pool to compensate.

- Length and complexity of the hold. A buy-and-build with a five-year hold and an expected dozen acquisitions will see more leadership turnover and more new senior hires than a stable carve-out with a three-year hold. The former needs room in the pool for future additions.
- Management's prior experience. A first-time CEO at the helm of a PE-backed business may need a larger relative allocation to be drawn in, they are taking real career risk. An experienced PE-CEO who has done this before may price their participation more sharply because they know the gearing.
- Sector benchmarks. Some sectors have conventions. Healthcare and consumer tend to run larger pools; industrials and infrastructure tend to run smaller. These conventions reflect both leverage levels and historical leadership intensity, and they should inform, not dictate, the answer.
- Tax treatment. What matters to management is the after-tax outcome, not the gross allocation. Tax treatment of MIP participation varies materially by jurisdiction, and within a jurisdiction, by the structural choices made at inception (employee vs investor characterisation, entry valuation, holding vehicle, timing of equity acquisition). Two programmes with identical headline pool sizes can produce very different net outcomes for participants. Sizing decisions should be informed by the expected tax treatment, not arrived at separately from it. This is one area where early specialist input, both tax and structural, can have outsized impact on whether the programme actually motivates.

None of these variables produces a single number. Together they produce a defensible range, and the act of working through them is itself the value of doing sizing properly rather than picking a percentage off a slide deck.

An honest opinion: most programmes are underdimensioned

Across the deals we see, the pattern is consistent. MIP pools are sized at closing with the participants present at closing in mind. They are not sized with five years of holding-period dynamics in mind. The result is that pools tend to be too small, not always dramatically, but systematically.

There are two reasons for this. One is about psychology at closing. The other is about mechanics over the hold.

REASON ONE, CLOSING-DAY ASYMMETRY

At closing, dilution is concrete and motivation is hypothetical. A pool of 10% of ordinary shares versus 13% is a tangible €4–5m difference for the PE fund manager at a strong exit, visible in the model. What you cannot see in the model is the upside that a more strongly aligned management team might have delivered with a larger pool. Maybe a €260m exit becomes a €310m exit; the PE fund manager would be materially better off net. But that is hypothetical. The dilution is real and visible.

Negotiation dynamics compound the bias. Management teams at closing are typically in a weaker position than the PE fund manager. They want to close the deal and are reluctant to push hard on the size of their own allocation, it feels like greed. Funds, consciously or not, take advantage of this. The pool settles at a number that is slightly smaller than optimal, and nobody on either side feels the cost until much later.

REASON TWO, THE POOL ERODES OVER THE HOLDING PERIOD

This is the stronger argument, and it is mechanical rather than psychological. A pool sized at 10% at closing is rarely still 10% in economic substance five years later. Three forces work against it:

- Additional equity rounds. Buy-and-builds, large operational investments, and refinancings can require additional equity to be raised mid-hold. If management does not participate *pari passu* (or cannot, or is not invited to), their proportional position dilutes.
- Good leavers retain shares. When a participant leaves as a good leaver, they typically keep their vested shares, or the company buys them back at market value. Either way, those shares no longer sit in the pool available for new allocations. The pool shrinks economically while the headline percentage may stay the same.
- New senior hires need allocations. Five-year holds rarely retain the original leadership team intact. A new CFO, a new COO, a new division head, each new senior hire needs a meaningful allocation to be motivated. That allocation has to come from somewhere.

The point is not that any one of these forces is catastrophic. The point is that they almost always run in the same direction, reducing what is effectively available for the team in place at exit. A pool that was already on the lean side at closing becomes uncomfortably tight by exit.

The right pool size at closing is not the size needed for today's team. It is the size needed for the team that will be in place at exit, five years and several leadership changes from now. Almost nobody sizes for that.

Mid-life additions: four ways to fund a new hire

Suppose two years into the hold, a new CFO needs to be brought in. The market expectation for that role might be a 1.5% allocation. Where does that 1.5% come from?

There are four practical approaches. Each has different consequences for the existing management team and for the PE fund manager. Most documentation does not specify which approach will be used, leaving it to be negotiated, often awkwardly, at the moment it matters.

METHOD	HOW IT WORKS	WHO BEARS THE DILUTION
A, From the existing pool	The new hire's allocation comes out of the original MIP pool. The pool shrinks for everyone already in it.	Existing management collectively. PE fund manager is not diluted.
B, Proportional issuance	New ordinary shares are issued and allocated to the new hire. PE fund manager and existing management both dilute proportionally to their current holdings.	Both parties, in their existing proportions.
C, Pre-reserved top-up	At closing, the pool is intentionally sized larger than needed for the initial team, so for example 12% rather than 10%, with the excess reserved for future hires and reallocations.	The dilution is priced in upfront, so neither party experiences a mid-life adjustment.
D, PE absorbs alone	New ordinary shares are issued only against the PE fund manager's holding. Existing management's percentage stays intact; PE absorbs the full dilution of the new allocation.	PE fund manager alone.

Each method has a logic. Method A is operationally simplest and is the default in many programmes; it also penalises the management team that is doing the work. Method B is the most symmetric. Method C is the most disciplined, it acknowledges in advance that the pool will need to flex. Method D is the most generous to management; it tends to apply when the PE fund manager believes the new hire is essential for delivering the exit and is willing to pay for them.

The strongest piece of advice we can offer here is the simplest one: decide which method will apply before you need it. Programmes that specify the mechanism in the original documentation handle mid-life additions cleanly. Programmes that leave it open get into uncomfortable conversations at the worst possible moment, typically when the business needs a new hire urgently and has limited negotiating leverage.

What this means in practice

FOR FUND MANAGERS

Size the pool for the team you will have at exit, not the team you have at closing. The pool will erode in ways that are difficult to predict but easy to anticipate in aggregate. Building in headroom, through Method C, or simply by sizing slightly more generously than your model suggests, is cheaper than renegotiating mid-hold.

Specify the dilution mechanism in advance. Documentation that addresses how new hires will be allocated equity prevents friction when the moment arrives. The right answer depends on the deal; what matters is that there is an answer.

Be honest about the asymmetry. The cost of a slightly larger pool is visible and quantifiable. The cost of a slightly smaller pool is invisible and shows up only in retrospect, in a softer exit or a team that disengaged earlier than you noticed.

FOR CFOS AND MANAGEMENT TEAMS

Understand what you are joining. If you are being asked to participate in a MIP or MEP, the label tells you almost nothing. What matters is the size of the pool, your allocation within it, the leverage in the capital structure above you, and the leaver provisions that govern your exit. Article one of this series covers the mechanics in detail.

Ask about the dilution mechanism for new hires. If you are joining an existing programme, find out how the pool has been structured to handle new joiners. If you are negotiating a programme at closing, raise the question explicitly. The answer will tell you a great deal about how the fund manager thinks about your team's interests over the hold.

This is the single question we recommend every CFO and incoming senior hire to ask, and almost nobody does. In most documentation it is genuinely unspecified, meaning the answer will be negotiated, on the fly, at the moment a new colleague needs to be brought in. By then, the leverage has shifted. Ask the question now, when there is still time to write down the answer.

Negotiate the structure, not the label. Whether the programme is called MIP, MEP, sweet equity, or management participation is irrelevant. The terms inside the documentation are everything.

One last thought

The right pool size is the one that works for the deal in front of you. Not the one your last deal used, not the one the industry benchmarks suggest, not the one your competitor disclosed in last week's case study. It is the size that aligns this team with this capital structure for this hold, and that requires a conversation, not a number.

The labels, MIP, MEP, sweet equity, are mostly interchangeable. The design choices behind them are everything. The question worth asking is not "is it a MIP or an MEP?" but "is the pool sized for the team that will deliver the exit, and is the mechanism for handling change written down?"

Most programmes do not get this right at closing. The good news is that it is fixable, mid-hold, with deliberate work. That is what we do at MIP Desk.