



Catalysis



What can we learn from 853 references?

Dr Mike Hicks
January 2015



Almost everyone involved in working with the management teams of growth companies gets involved in referencing incumbent (in a deal situation) or new (due to churn) executive directors. Most investors seem to regard referencing as an important tool but, in many cases, one which is approached inexactly and yields mixed results. The quality of references can suffer from deliberate boosting (or sabotaging) by referees, or a softer less conscious positive bias. Academic research on this topic is limited but what there is offers only lukewarm validation of the value of referencing. The problem for both practitioners and researchers, though, is that much referencing practice is conceived vaguely and executed inconsistently.

To throw some light on this topic, I have carried out analysis of reference calls I made for clients (>95% private equity related) over the last five years. After removing calls without full information, I ended up with 853 calls. As Chart 1 shows, 15% of these were references taken on whole teams/businesses as opposed to individuals – and are excluded from the analysis on subsequent pages.

GRADING AND SCORING

Reference calls – if given the time they deserve - are complex things potentially involving relationship context; facts about previous roles; achievements, strengths & weaknesses; views on suitability; advice on working with the candidate. Reducing all of that to a single grade is inevitably an imperfect exercise in judgement and approximation. The descriptions of the grades I use can be found in Table 1, along with the scores I use to create weighted averages and benchmarks. The actual names don't matter as such, but the differences between them do.

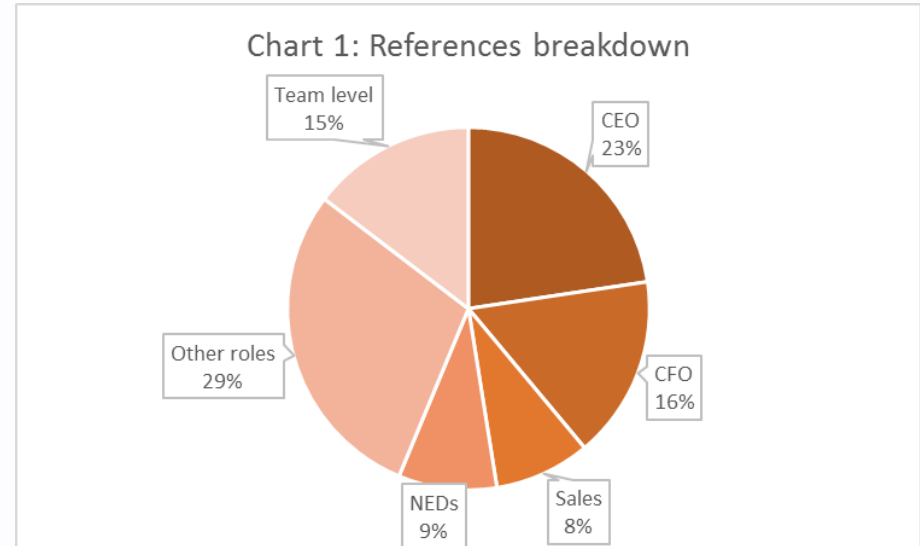


TABLE 1

Grade	Description	Score
Very high	Consistent and substantiated evidence of high performance relevant to the role.	+5
High to very high	Substantiated evidence of performance, but some inconsistency or other issues.	+3
High	Generally positive, but lack of substantiation or more than one issue present.	+1
Medium high	Some positives, but important gaps and/or significant unhelpful issues.	-1
Low/medium	Real question marks about suitability of the candidate.	-3


Referencing is often seen as a way to check whether there are negatives about a candidate. Some exercises can indeed become forensic examinations of fraud, bullying, volatility and gross incompetence: I have some great stories in each of these categories and more.

But the mundane reality is that most candidates are not hiding dark secrets – only lower levels of human frailty. Chart 2 shows that the two bottom categories ('Medium high' and 'Low Medium') represent just 10% of the calls made. Of course, I may simply have missed bad stuff because referees were coy or I didn't spot their hints. However, any such gaps appear limited in number and impact.

Instead, the difficult part about referencing is distinguishing between the three larger and more positive categories ('High', 'High to Very High', 'Very High'). For every bad 'un I reckon there are five who turn out to be mediocre, disappointing or underwhelming. This implies that it is a lack of substantiated and consistent positives which is the main cause of failure.

That begs the question: which positives are actually needed in the current/proposed role? If we know what the role requires then we can spot not just which strengths are validated but which ones are not mentioned at all. Those 'missing positives' may be more likely to point towards failure than various explicit negatives. The latter can often be steered around whereas gaps which have not been identified cannot easily be mitigated.

Chart 2: Reference grades - all individuals



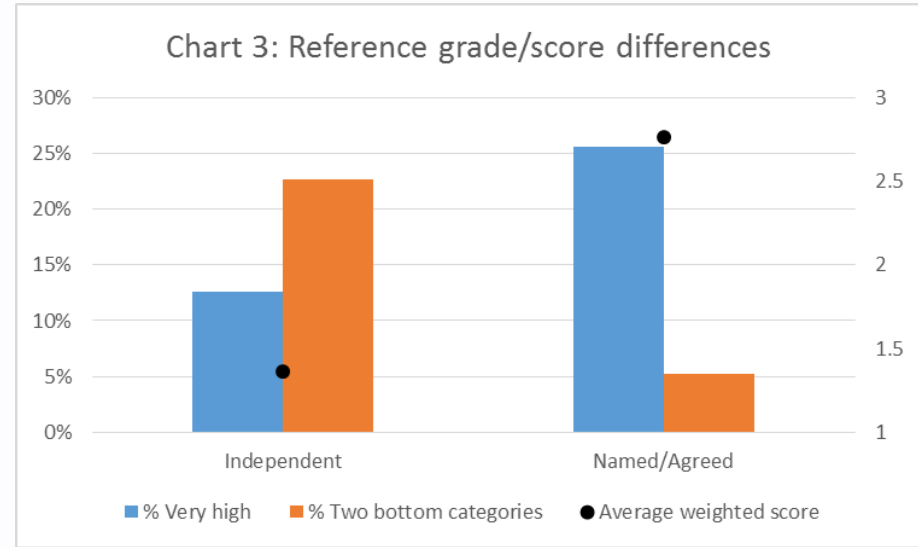
What can we learn from 853 references?

It makes intuitive sense that named referees are more likely to speak positively of a candidate than an independent referee. In general, that is true but it is worth noting that:

- ▶ Some candidates show sufficient lack of judgement about the situation and their contacts that their nominated referees either reveal serious issues – even while insisting what a good fellow they are – or actually explicitly talk them down. As Chart 3 shows, about 5% of named referees do this so, personally, I will almost always take up referees even where they appear to be as much friendly contacts as genuine professional contacts. You never know what they might come up with! In practice, many more than this produce middling ‘High’ scores where positivity is expressed but the referee struggles to articulate actual strengths.
- ▶ Plenty of independent referees are convincing advocates for candidates. Many executives assume that past colleagues have forgotten them or might have little to say: in practice, referees are more vocal and forgiving than worried candidates may anticipate.

Nonetheless, the big message of Chart 3 is that independent referees are indeed less likely to offer top grades to candidates and more than four times more likely than named/agreed referees* to offer ‘Medium High’ or ‘Low/Medium’ grades. Converting that into a weighted score, nominated referees are about almost 1.5 points higher in their scores than their independent counterparts.

How much ‘softer’ are named referees?



*I refer in the chart to ‘agreed’ as well as named referees. Typically these are the result of a conversation with a candidate, often after interview, where names are discussed, some eliminated and others added. They are contacted by candidates in the same way as named referees, but the relevance of the referee set has been boosted.

What can we learn from 853 references?


Referees are not all created equal. As Chart 4 shows, some referee types are more likely to express positives and less likely to mention negatives.

- ▶ Subordinates are the most net positive, although independent referee subordinates have sometimes provided the most damning accounts of certain individuals.
- ▶ Referees who have limited professional involvement with the candidate, or paid advisers, (referred to here as 'friend etc.') are next.
- ▶ NEDs can be both amongst the most positive but also more critical than average.
- ▶ Industry observers, clients, suppliers (called 'Industry etc.' here) and more senior colleagues are more middling.
- ▶ The most critical, by a clear margin, are investors. So the instinct to call a PE colleague who has previously backed a candidate is a sound one if you are looking for tough realities.

Although some referee types bring higher levels of scepticism, talking to all types is useful since it is often through an accumulation of commentary describing similar attributes which helps make clear what is actually meant. So the key criterion for referee selection should be potential insightfulness rather than hierarchical position. Likewise, especially in the case of incumbent teams where members have been in the same role for more than five years, it is often former employees or clients who can provide recent views which, for example, bosses from years ago cannot.

Which referees are the hardest to impress?

Chart 4: Reference scores by referee type



What can we learn from 853 references?

I am not keen on taking references from referees who have not worked with candidates in the last ten years. One I did do described an investee CEO as a 'nice young man' (he is about 50!). But my analysis (not shown here) shows no link between time elapsed and overall positivity.

By contrast, it appears that:

- ▶ The closer the referee worked with the candidate, the more likely they are to provide the kind of detailed substantiation of strengths which allows 'Very High' grades (see Chart 5). By contrast, referees who worked less closely may give a more approximate sense of the candidate where that is more difficult. But distance also seems to generate a higher degree of critical commentary: see the unusually high levels of low grades from the 'Not close' group.

- ▶ The longer referees have interacted professionally with candidates, the more positive – and less negative – they tend to become (see Chart 6).

Without doing some heavy-duty statistical analysis it is hard to disentangle the relative effects of referee type, closeness and time worked together. But intuitively, we might expect that those who have worked together longest may be, for example, long serving subordinates or colleagues who have worked closely and who are offered as named referees. By contrast, investors and NEDs are likelier to have shorter periods of interaction.

What is the effect of time and distance on grades?

Chart 5: Grades by closeness of relationship

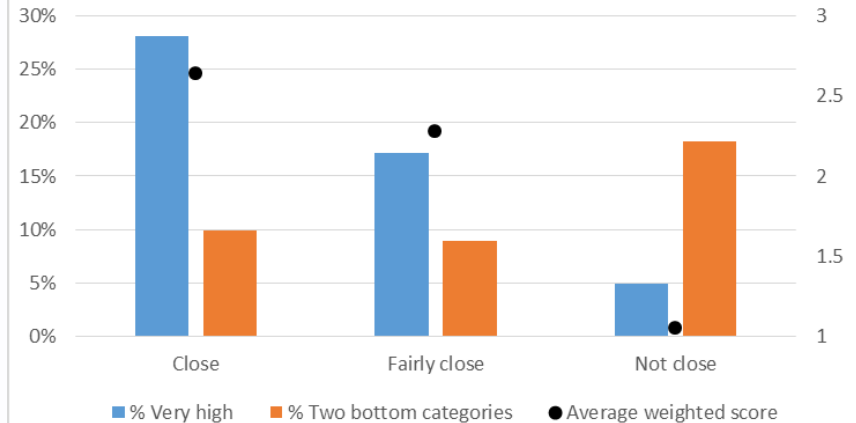



Chart 6: Grades by length of relationship



What can we learn from 853 references?

Who do referees most rate?

Charts 7 and 8 suggest that not all candidate types attract equal support:

- ▶ Individual executives already in post are less likely to attract the highest ratings than those introduced new to a business. This makes sense since (a) teams often stay constant even as business needs evolve away from the internal skill set – especially as Chief Executives/boards may only remove directorial colleagues for serious under-performance; (b) new candidates have typically been selected from a range of candidates keeping the most recent requirements in mind. Moreover, I tend to reference finalist candidates who have survived initial rounds of interviews which tend to remove obviously weaker people. But note that it is not more negatives which distinguishes between the in situ and new candidates, rather a lower proportion of really strong grades for the in situ people. This supports the contention on page 3 that what typically distinguishes between candidates is the presence or not of substantiated positives rather than the identification of negatives.
- ▶ As often seems to be the case in private equity situations, FDs/CFOs seem to get a harder time than other candidates! They are less likely to attract positive grades than CEOs, and more likely to attract negatives – although the differences are not huge. Curiously, NEDs (including chairmen) seem to generate more variable feedback: they attract more 'Very high' grades but also a somewhat higher proportion of criticisms.

Chart 7: Scores by status of candidates

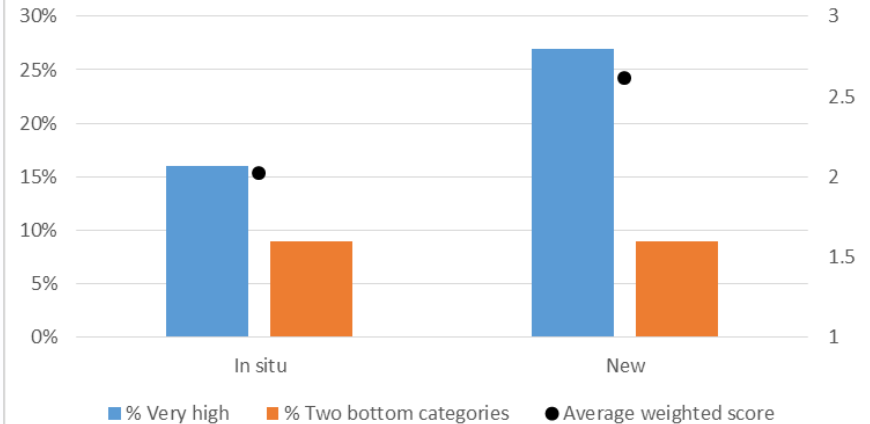



Chart 8: Scores by proposed role



What can we learn from 853 references?

Based partly on the analysis above, but also almost ten years of referencing candidates and teams for PE houses, my top tips are as follows:

- ▶ As with any aspect of assessment, working out what you actually need has to be the starting point. What is critical and what is just a preference?
- ▶ Referee numbers matter if you hope to get a minimally comprehensive view. Take at least 6-7 for any directorial position; more like 10 for a Chief Executive role.
- ▶ Referee selection matters, so agreeing names with candidates to ensure they are relevant to what you want to find out is important.
- ▶ Independent referees are often more revealing than named/agreed ones. But it is only polite to let candidates know you intend to talk to them and let them have a chance to explain if there is anyone specific it would be a real problem to contact (and why). Proceed within agreed parameters but keep their names anonymous (although some will almost certainly tell the candidate – hence the need to discuss the principle in advance).
- ▶ Identifying independent referees is much easier than in the past using the web, company house data and tools like Linked-In. If you don't have time to do this yourself then get someone else to do the legwork.
- ▶ I get better responses from named and independent referees when I write to introduce myself, and what I am doing, in advance. Calling out of the blue can make referees awkward and unwilling to open up.
- ▶ Unless the referee has had only marginal contact, a good call should last anything between 15 and 30 minutes. Having a general checklist of what you want to cover can keep the conversation going. Get referees talking in their own words about the background and relationship first before asking more pointed questions.

So how do we achieve better referencing results?

- ▶ Ask for specifics and examples – especially where language can be ambiguous. For example, 'does not suffer fools lightly' can cover everything from sensible handling of under-performance, through destructive perfectionism, to deranged angry men.
- ▶ If the referee (or you) find discussing negatives directly awkward then asking indirectly often opens things up. For example, 'What is your advice on working with Fred?' or 'What colleagues do they need to complement their pattern of strengths?'
- ▶ To gain perspective, ask referees for comparisons between the candidate and other people they have worked with in similar roles.
- ▶ Remember that individual performance can vary widely by organisational context and particular business challenges. So mapping referencing data isn't just a synthesis of pluses and minuses but rather an imagining of how knowledge, skills and attitudes will transfer.
- ▶ Whether formally or informally, what you hear from referees needs to be calibrated by:
 - Comparisons with other people in that role. I use benchmarks from previous projects, but even talking it through can be informative.
 - Thinking about the referee mix and its likely impact on positivity. Substantiated praise from independent referees (especially investors) means more than feedback from 'softer' sources.
- ▶ Performance is usually a function of an overall team so many individual issues can be worked round by tweaking roles and shuffling deliverables. That assumes there are sufficient strengths to be worth making that accommodation.



Catalysis

