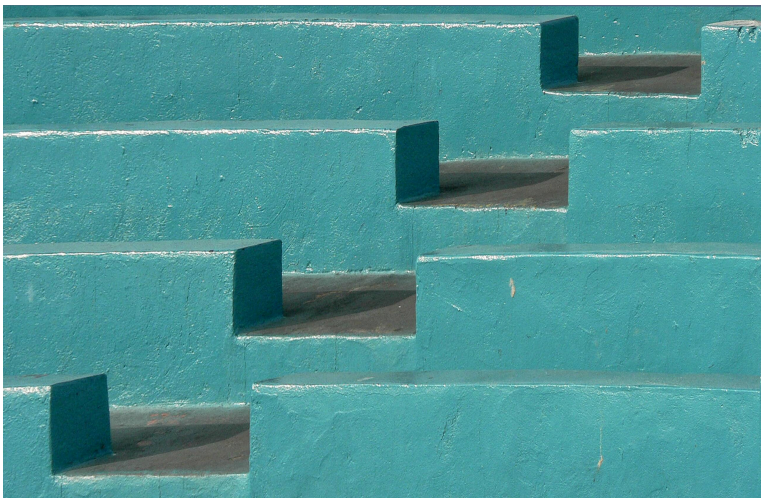
How Small, Fast Steps Can Drive Impact at Scale

Published on Apr 13, 2021 by Niko Canner



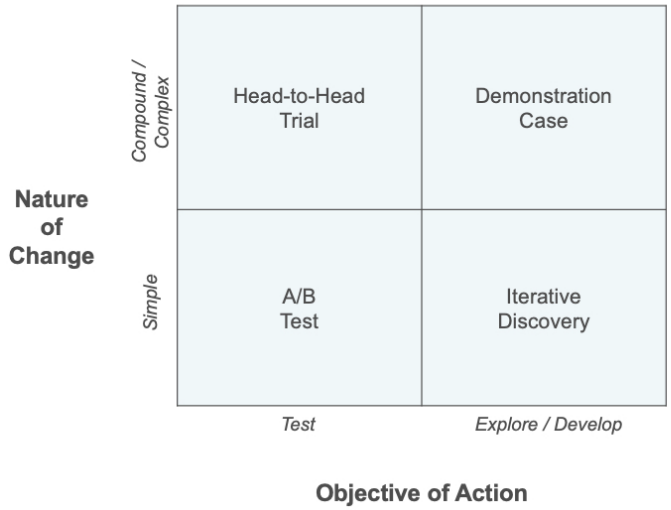
In an environment of rapid change, only those who adapt can succeed. Market leaders often find that the very same strengths that have fueled their historical success—scale, breadth, and the ability to knit together complex processes—hold them back. Traps abound. Move too slowly: others who adapt faster will gain share. Make overconfident big moves: too many will be wrong. Experiment effectively, but don't scale: proliferate complexity and cost without moving the needle. What's needed is a discipline to take small, fast steps to push beyond the current envelope of performance, learn systematically, and scale what's demonstrated to work.

Part of what's challenging about developing this discipline is that, peeling back underneath the headline point, there are very different ways to think about the small, fast steps depending on the objective of the action and the nature of the change. Leaders find a "piloting approach" that's worked for them, and risk applying that too broadly, versus fitting the approach to the need.

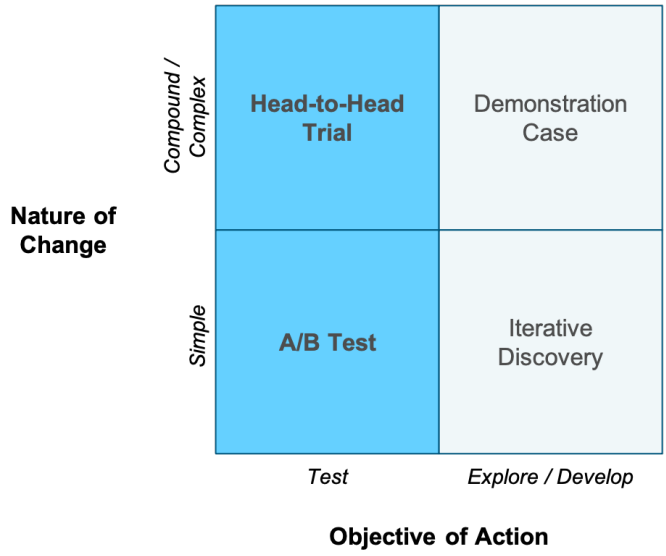
Designing to fit the need begins with answering two questions:

1. What is the objective of action: are we testing a model we can fully specify up front, or are we exploring and developing a model in action?
2. What is the nature of the change: is it a simple, discrete shift from what we do today that can be isolated; or is it a complex, compound change made up of many interdependent moving parts?

These two dimensions point to four distinct disciplines to taking small, fast steps designed to pave the way to action at scale.



Action to Test, Validate and Refine



On the bottom left, the simplest testing action, is the discipline of the **A/B Test**. Here, we're in the domain of shifts to micro elements. Should the text be blue or gray? Should the rep begin the call with this line or that one? The whole domain of digital marketing is built on these disciplines, and a great deal has been written about the associated tools and management approaches. At the heart of managing this discipline well is throughput volume of running lots and lots of tests, building a system to generate variation, isolate what works, and continuously refine.

Most organizations have internalized this discipline in certain parts of their operations, such as digital marketing and direct mail. Whenever this discipline can be extended to a new domain, there

can be extraordinary returns. More than twenty years ago, my team and I were working with the credit card collections division in one of the large banks, and one second-level supervisor in Tampa, a woman named Charlotte Fludd, was outperforming her peers so much that the delta of her unit's performance was a million dollars in value per month. The essence of Charlotte's outperformance was that while other supervisors were changing many things at once (what customers they targeted, when, with what approaches), Charlotte worked by the scientific method, shifting only one element at a time in her team's operational calendar, so that she could isolate which changes improved performance. That's the A/B testing discipline in action. (Charles Duhigg tells this story in greater detail in his book Smarter Faster Better.)

Many changes, however, can't be reduced to independent, discrete elements. For instance, consider a shift in the selling model that involves a new framing of the company's value proposition, new targeting models, new sales aids and retraining of reps. This model can't be atomized into hundreds of discrete moves to be A/B tested. The elements of the new model hold together as one interdependent whole.

A company wouldn't undertake the investment in building the new sales model if there wasn't a basis for conviction regarding the performance lift from making the change. At the same time, it would be foolish simply to assume that the improvement will manifest, and to roll out nationally (or on an even broader scale) without some kind of test. The relevant discipline here is the **Head-to-Head Trial**: setting up a fair, contained test of the new model, implemented as a whole against the relevant baseline.

Effectiveness at this discipline hinges on quality of measurement, which comes down to three factors:

- Rigorous comparison of experiment and control
- Avoiding contamination
- Attribution of causality for observed results

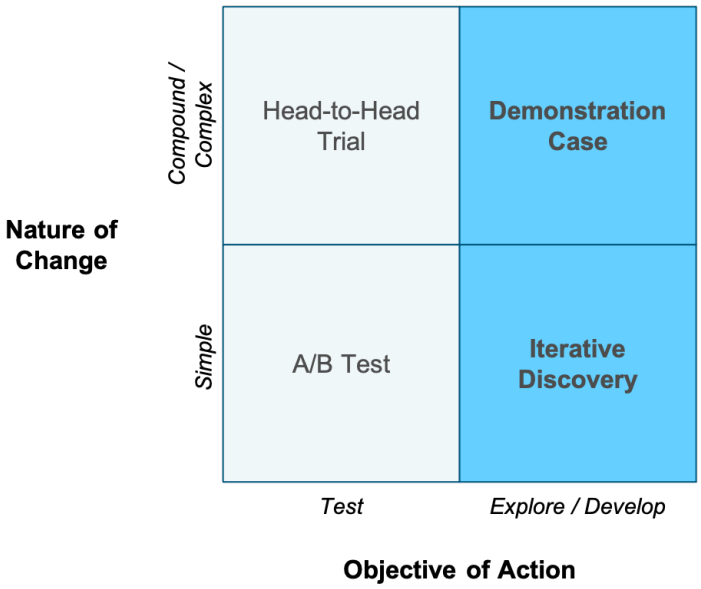
In a scientific experiment, the experimental group and the control group are otherwise identical and differ only in terms of the relevant "treatment." This is difficult to achieve in business. For instance, suppose a new sales model is tested in Region X at time T. The results can be compared to the rest of the country, but there may be other differences between Region X and others, separate from the experiment (e.g., different competitive conditions). The results can be compared to Region X's own results in a prior period, but there may have been other shifts that happened between T and T—1, which have nothing to do with the experiment. While head-to-head trials in business are almost never as pure as scientific experiments, they can be designed carefully to maximize accuracy and minimize contamination. One frequent and avoidable source of contamination is executive focus. Because executives (rightly!) care so much about the result, they pay attention in all kinds of ways that won't be scaled when the pilot is in full roll-out (e.g., they surface and address barriers while the work is in progress, they assign higher-performing managers into the experimental region, they shift what isn't working mid-stream). This can produce a "Wizard of Oz" pilot, in which the "man behind the curtain" is redesigning the experiment even as a head-to-head trial is meant to be occurring.

The results of head-to-head trials can be murky, and even when results are clear, the challenge of scaling up a new model with many interdependent moving parts can be significant. Mastering these dynamics requires thinking up front about instrumentation and learning loops, so that, if there's a surprising result (e.g., the new model doesn't drive the level of lift expected), we can decide whether there's reason to believe:

- There's something wrong with the fundamental premise of the model
- Some particular part of the model isn't working in the way it is intended (e.g., the value proposition is right, but the sales aids don't convey the value proposition powerfully)
- The model is strong, but there's an issue with the quality of implementation (e.g., some reps have the motivation and the skill to make the transition and others don't)

The art and science of the learning approaches involved here could be their own treatise. As shorthand here, the disciplines involved are the opposite of how many managers intuitively operate: their discipline is to perceive what they think isn't working and immediately take action to fix it; where the discipline of the head-to-head trial is to focus on inquiry to understand what's driving variance, while leaving the trial itself untouched in order to ensure there's a valid comparison. By maximizing learning during the trial, not only can one make the most valid assessment of how well the model worked, one can understand the barriers to effective implementation at scale that need to be addressed in order to move forward. Much of the relevant learning on this point is qualitative in nature, understanding where the team members involved encountered issues they needed to address, mixed messages, and tensions between what they perceived would be effective and how the model they were implementing had been designed to work.

Action to Drive Discovery and Develop New Models



In the section above, we've been focused on applying a known model at small scale, in order to measure results, test and scale what's been demonstrated to work. A different set of disciplines apply earlier, when there isn't yet a fully working model, but rather a hypothesis or a strategic concept that needs to be fleshed out and manifested.

In the lower-right, the simpler explore/develop action, the purpose of taking small, fast steps is to drive **Iterative Discovery**. Rather than learning by studying a problem, the discipline here is to take action in small ways that generate insights, which over time can be synthesized into a strategy, a policy or an operational process.

For example, in the early 2000s, Pfizer, already the "powerhouse marketer" in their industry, launched a Marketing Innovation team to push the envelope in how they engaged customers and created value beyond the pill. Rather than think big thoughts and write PowerPoint presentations, Ellen Brett—who had been appointed to lead this team and who later became Chief Marketing Officer of Pfizer's $20B Primary Care business—and I built a small team with counterparts in Sales and Medical. We decided that rather than get distracted by the complexity and variety of the healthcare ecosystems in which Pfizer operated, we'd be better served to learn intensely in one microcosm and then see how this learning generalized. Thus began the "Allergist Experiment," which started with identifying eight allergists with different kinds of practices, understanding their needs and aspirations deeply, and asking ourselves: "if all of Pfizer's expertise and capability could be put in service of helping these allergists improve their clinical outcomes and advance their aspirations for their practices, what might be possible?" Of course, there are many constraints regarding how pharmaceutical companies can engage with physicians, which limited what ideas we could act upon in what ways, but we were careful not to let those constraints limit our insight into the underlying realities of our customers. We wanted to ensure we saw the world through their eyes and then brought Pfizer's knowledge and capabilities to bear, rather than trying to bring them around to seeing the world in our way. We prototyped ideas in hours or days, making sure in the early going that we didn't worry too much about what would scale, and drove fast feedback loops. The sparks that flew from the allergist experiment became foundations for years of work to evolve how Pfizer thought about physician insight and about how to position the company and not just the brand, which were applied across a very wide range of physicians.

With the Allergist Experiment, our goals were initially defined at a very high level. We deliberately left ourselves open to a wide range of possibilities, so that we'd learn through creative engagement with physicians rather than let ourselves be limited by the opportunities we could see at the experiment's beginning. Moving to the upper-right box, a complex discovery action, the goal is more sharply and tightly defined, but the recipe for achieving the goal isn't yet known. In this **Demonstration Case** discipline, the purpose of small, fast steps is to build toward practical understanding of how to deliver on a known goal and to demonstrate that the goal can be achieved.

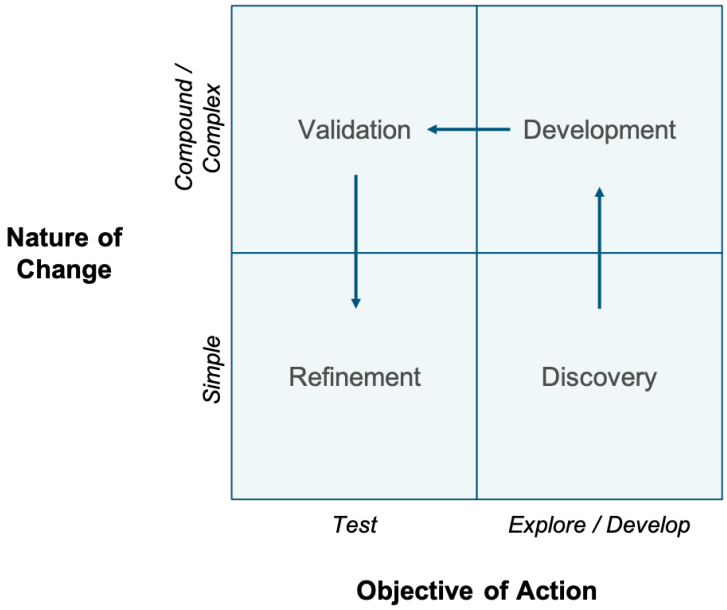
For example, working in a call center focused on inside sales many years ago, we quantified that first-level supervisors were spending only 15% of their time coaching reps—and much of the rest of the time on various forms of administration and reporting. If that 15% could be quadrupled—so that supervisors functioned primarily as coaches, with 60% of their time helping reps progress toward the next level of effectiveness—that would create tremendous lift in the customer experience and in sales outcomes. As we began talking about this opportunity, there was universal agreement that this

would be an immensely valuable shift. However, there was widespread skepticism that such a large shift would be possible. Supervisors felt that there were many internal barriers to making this kind of change, and there wasn't anyone whose job it was to solve for the many moving parts of eliminating, simplifying, automating or reassigning the many tasks that made up the other 85% of the supervisors' time. Change couldn't be dictated from the center—after all, each supervisor made his or her own decisions about what to do throughout the day, rather than following a set operational schedule shaped by headquarters.

To break through his impasse, we worked closely with one site, in Tempe, Arizona, to create a demonstration case. Our team worked closely with a small group of supervisors to clear away as much of their non-coaching work as possible, troubleshooting how to solve each of the operational problems we encountered along the way. This pilot had support not only of site leadership, but all the way up to the national leader of Sales, so that as we encountered issues related to policy and systems, we had aircover to work with the necessary people to find some other way that the relevant needs could be met (or could be set aside), so that supervisors could focus on coaching. The supervisors tracked their time rigorously, and week by week, we assessed our progress toward the 60% goal, like a funding campaign tracking their progress on a giant thermometer. After several weeks of work, the goal had visibly been achieved. The supervisors in this particular unit in Tempe became ambassadors to their peers about what was possible, and we could build on what we'd learned to distill what we'd built from the ground up into a scalable, teachable approach.

As the supervisor coaching story illustrates, the discipline of the Demonstration Case is in many ways the opposite of the discipline of the Head-to-Head Trial in the upper-left. In a Head-to-Head trial, one focuses on a fair test of the new model against the comparison. To achieve this, the new model needs to be pinned down in advance, and the "Wizard of Oz" trap must be avoided. In the Demonstration Case, the whole point of the effort is to drive work "behind the curtain" in order to solve a problem that hasn't been solved before the work begins. It's almost never true that an operating unit has the capacity to drive all this invention and transformation work with its existing resources in a way that would represent a fair test. The operating unit needs the extras of executive sponsorship, budget, staff and/or consulting resources, and so on in order to figure out how to achieve the goal, solving the problem in action rather than solving the problem on paper.

In practice, there can often be a natural flow among these four disciplines of small, fast steps as an organization progresses from high-level hypotheses, to concrete stretch goals, to testable models and ultimately to a working design at scale that needs ongoing refinement and continuous improvement.



This arc begins on the lower right, driving Discovery through rapid, iterative action that generates learning. This learning cements conviction about the right stretch goal to be achieved—but there is still not a model that establishes that the goal can be achieved. The demonstration case discipline accelerates Development of a working model. This model still needs Validation, as the process of the demonstration case can't be replicated at scale. The model is distilled into a formalized, replicable, teachable protocol, which can then be tested through a Head-to-Head Trial. If that trial is successful, implementation at scale proceeds. Even when implementation succeeds, there's still a need to drive continuous improvement, which now proceeds not in the big leap of a transition from one model to another, but through careful, frequent, disciplined Refinement, isolating specific changes and testing their effect.

Knowing why one is piloting, and designing the pilot accordingly, is what makes these practices a set of distinct disciplines. After all, developing a model that demonstrates what is possible (Development of Demonstration Case) requires a very different discipline than testing whether a model works under certain conditions (Validation in a Head-to-Head Trial). In selecting the right discipline, leaders can push beyond their current performance in a directed way—using small, fast steps to learn systematically and scale what's demonstrated to work.

Photo by Marco Bianchetti on Unsplash