

An Introduction to Evidence-Based Portfolio Management

By [Scrum.org](https://www.scrum.org)

Organizations seek agility to improve their competitiveness by being more responsive to change. While many organizations have reaped the rewards of agility at the team level, they struggle to reap these benefits more broadly because their traditional management practices stand in the way. Agile principles and practices must spread beyond the Scrum Team in order for organizations to achieve the dramatic improvement that they seek in their business results.

Clearing the way for greater success requires organizations to change the way they invest in innovation and artificial and quantify success. Organizations that think they simply need to deliver faster, without considering what they are delivering, are missing the greatest benefit of agility: rapidly trying out ideas, discarding those that don't work, enhancing those that do, and uncovering new ideas from unexpected sources. Being able to change investment strategies, quickly, based on new information, separates the companies with winning products and happy customers from those companies who are mere footnotes in the history of their industry.

Traditional Portfolio Management Wastes Time, Effort, and Money

Traditional portfolio management focuses mostly on activities and outputs, and only somewhat on outcomes. At the front end, there is a massive intake process in which a bunch of definitions of solutions are collected. This process is frequently restricted to an annual budgeting process, so ideas that occur at other points in the year are often lost. The outcomes that the solution will achieve are often only vaguely defined.

People who deliver solutions (usually their managers, not the people who will do the work) are asked to estimate the cost of building the solution (resulting in a rough plan of activities, including a rough schedule). While these estimates have all manner of caveats wrapped around them, the caveats are typically ignored and the estimates and rough schedule are enshrined in the portfolio management system, where they become weaponized as hard, non-negotiable commitments by I.T. to Business.

When funding decisions are based on unrealistically optimistic estimates made by people far removed from the actual work, the size of the work later expands when the real workers get involved. This increases the price tag of the solution, causing sticker-shock in the Business organization, who feel like victims of a bait and switch. This further corrodes the bond between Business and I.T.

Sometimes organizations manage this risk by engaging some representatives of people who do the work in the big-bang portfolio estimation exercise. While this improves the quality of estimates, it also distracts developers from delivering work that is already committed, which also erodes the delivery confidence of their business partners. Whether they engage or not, developers are saddled with unachievable delivery expectations.

The PMO tries to balance the benefits and costs of the solution based on guesses about a few key variables (cost, time, importance, governance, etc.). Then, politics enters the process, priorities are negotiated, some “horse trading” results in some shifts here and there, and then the budgets for the year are more or less locked-in. When new and better opportunities arise, the change management overhead is so intense that rapid response to new opportunities is nearly impossible unless the sponsor is willing to spend a lot of political capital.

Sometimes the situation is opposite. When the size of the projects are huge, the risk of failure can be paralyzing, leaving budgets that are always “in discussion”. The opportunity eventually slips away because the organization required a high-fidelity plan that could meet all needs and address all risks, before it would invest the money to proceed.

In either case, the traditional approach prevents the organization from making the best of its opportunities:

- The “solutions” are really just conjectures (sometimes referred to as HiPPOs, or *Highly Paid Person’s Opinions*) about what might be needed to deliver one or more outcomes
- The specification of the solution is vague and imprecise, and when it is not, it is bloated with features that don’t address real needs or improve desired outcomes
- The refinement of ideas often fails to quantify the hypothesis and build a forecast of expected measurable value. The item rarely includes the means to validate the hypothesis by comparing expected measurable value to actual measurable value
- Cost and schedule estimates are wild guesses based on minimal information, and are usually produced under extreme and artificial time pressure during an annual budgeting process
- The people producing the estimates (usually managers) are too far removed from the real work to produce a reasonable estimate, even if the solution desired was correct
- Having managers make commitments on behalf of teams damages the teams’ ability to self-organize and be accountable for their results
- Holding teams accountable for delivering to wildly flawed cost and schedule commitments is demoralizing and demotivating
- The “solution” often doesn’t deliver on the vague or unspecified goals that were originally envisioned
- Failure to validate the original hypothesis to extract learning dooms organizations to repeat the mistakes of the past in future portfolio management cycles
- Good ideas come up all the time; locking into a once-a-year plan can prevent the organization from seizing new opportunities

A more certain process for wasting time, effort, and money, and destroying value, morale and reputations, is hard to imagine. It's not the fault of the people creating the estimates or the teams delivering the solutions; they are trapped in a fundamentally flawed system that makes good people do bad things to each other. Also, governing projects against schedule, budget, and output commitments made during this flawed process compounds the problem.

Focusing on Outcomes Produces Better Results

Organizations fund work to achieve something, some measurable *outcome*, usually expressed in terms of some improvement in a customer experience. Outcomes are often discussed in terms of a customer's *jobs to be done*, or things that they need to, or would like to accomplish.¹ When these outcomes are realized, they produce *impacts*, which are things that benefit the organization like increased market share, revenue, or profit.²

Rather than funding activities that produce outputs of questionable value, be open and direct about the goals that the organization needs to achieve, and then let the people closest to the work determine how to achieve those goals. Centralize the organizational mission, vision, outcomes, and strategy and decentralize the Product Vision, strategy, and execution. Use frequent inspection of market-based evidence to govern the investment. In short, set clear and aspirational goals to aid alignment and commitment, and let the people closest to the work and the customer develop creative solutions that create real value. We refer to this as *Evidence-Based Portfolio Management*.

What is Evidence-Based Portfolio Management?

It is an approach that applies lean and agile principles to the challenge of deciding where to invest to derive the greatest business benefit by using rapid delivery of small increments of value to gather information and adapt investment decisions based on evidence from the market. It enables organizations to quickly test ideas by actually building and validating the smallest solution that will deliver a single outcome to a single set of customers or users.

It avoids the typical bloated upfront over-specification experienced in traditional projects while taking solutions to market as quickly as possible. It also enables organizations to experiment, minimizing affordable loss by eliminating bad ideas early, and finding good, nurturing ideas by obtaining real-world feedback. It makes sure that the most viable ideas get the support they deserve, and that agile teams can focus on delivering great solutions. In place of expensive and ineffective project status reporting meetings, it uses direct evidence to inspect and adapt on strategic investments.

Evidence-Based Portfolio Management builds upon a related empirically-based continuous improvement approach defined by Scrum.org called Evidence-Based Management (EBM). Like EBM, It is inspired by practices that originated in the medical community in response to the need to improve treatment outcomes “by integrating the best available evidence with practitioner expertise and other

¹ <https://hbr.org/2016/09/know-your-customers-jobs-to-be-done>

² <https://www.scrum.org/resources/evidence-based-management>

resources, and with the characteristics, state, needs, values and preferences of those who will be affected".³

Evidence-Based Portfolio Management Takes a *Principles-Based* Approach

Every organization's structure, roles, processes, and responsibilities are different. Rather than superimposing new roles and processes on an existing organization, Evidence-Based Portfolio Management takes the position that making and monitoring investments is better described by a set of principles that define and constrain how the organization identifies opportunities, decides which opportunities to pursue, and uses experimentation to guide whether to increase, continue, or stop those investments.

Different organizations have different structures for achieving this work, and the structure, roles, and events are not as important as the philosophy that guides the decisions. The underlying principles of Evidence-Based Portfolio Management principles are:

1. Separate capacity-for-growth from focus-of-work
2. Make the best decision you can, based on the best evidence available
3. Invest in improving business impacts using hypotheses, don't just fund activity
4. Continuously (re)evaluate and (re)order opportunities
5. Minimize avoidable loss
6. Let teams pull work as they have capacity
7. Improve status reporting with increased engagement and transparency

Each of these principles is described in the next section.

Principles of Evidence-Based Portfolio Management

1. Separate Budgeting for Capacity from Investing for Innovation

Teams are the engines of work in a modern organization, so if you take on new work, you either need to add more teams or enable the teams to be more effective. Splitting a team's focus between different things actually makes them less effective, so adding major new initiatives means adding new teams. Executives decide to add more teams when they feel the teams will deliver sufficient value to justify the expenditure. What these teams will work on will change, over time, based on business needs and competitive pressures.

Teams take time to form; team members need to be hired or moved from other teams, and they take time to learn how to work together. Because of this, you have to anticipate demand, and only when you are sure that you have more work to do will you create another team.

You will always have more ideas for things you could do than you will have teams to work on them; The art of portfolio management is deciding what to *not* work on at the moment. The number of teams you have will limit how many ideas you can work on at once.

³ <https://www.scrum.org/resources/evidence-based-management>

The number of teams an organization has is determined by its budgeting process. What those teams work on is governed by the portfolio management process. Practices like *Beyond Budgeting* help organizations to simplify the process by which they make decisions about capacity.⁴ The rest of this paper will focus on the process of deciding how you should use that capacity to get work done to create the best overall business outcomes.

2. Make the Best Decision You Can, Based on the Best Evidence Available

The evidence on which you will have to make decisions is always incomplete and sometimes unreliable. An empirical approach accepts that evidence is incomplete and allows for it being unreliable because it always tests assumptions and seeks better evidence by taking small steps, measuring, inspecting the results, and then adapting as necessary.

Some evidence is more compelling than others. So, while the opinions of stakeholders or subject-matter experts may be the best one can do when an idea is proposed, organizations should seek evidence that can confirm, deny, or evolve these opinions. Case studies from other organizations may help, but the best data will come from purposefully designed experiments to test specific assertions.

Organizations should make their decisions based on the best evidence available. The amount of money they “bet” on an investment should depend on their degree of certainty about the viability of the alternative. When information is less reliable or the risk of a hypothesis being wrong is higher, organizations should “risk” as little as possible to improve the evidence that they have to make decisions. As the quality of the evidence and the viability of the alternative improve, they can increase the investment in the potential solution accordingly.

3. Invest in Improving Business Impacts Using Hypotheses; Don't Just Fund Activity

Traditional portfolio management is principally concerned with three variables: cost, schedule, and output. These three variables, unfortunately, have little connection to value: cost has a loose connection to value in that it might cost too much to deliver the value that customers desire, so high cost may make certain solutions undesirable; a schedule's connection to value is also tenuous, as when an opportunity must be seized within a certain window of time or it becomes irrelevant. Output, or productivity, is only relevant when the right solution is being delivered.

What you really should be concerned with is the business impact of what you do. You do this by delivering products and services that help customers achieve better outcomes, within a window of opportunity, for a reasonable price, at an acceptable level of cost. Since the world is uncertain, you can't plan what customers need up front, such as what they are willing to pay, and how much it will cost. This is why you need to work iteratively and empirically in order to achieve superior business results.

You start by looking at the outcomes that customers want to experience in order for you to achieve the impacts that you desire. Unsatisfied customer outcomes represent market opportunities that you can target to achieve the improved results you are looking for. For every opportunity, you should be able to form a hypothesis about the value of pursuing the opportunity based on the number of people

⁴ <https://bbbt.org/>

who will benefit if you can deliver the desired outcome, and the value of delivering that outcome. These measures form the basis for the *Unrealized Value* of the opportunity.

Traditional approaches also define the size of potential opportunities. What is different about an empirical approach is that you measure the results of your experiments to determine whether your hypothesis is correct. If it is not correct, you may want to shift your focus to a different opportunity. Even when it is correct, you may find that the benefit of achieving the outcome is not worth what you will spend to achieve it, and so you also shift your focus to other opportunities.

For example, when profitability matters, as it nearly always does, the organization should establish specific goals on outcomes to say, “if you think you will spend more than X to achieve this outcome, let’s talk because we probably don’t want to pursue it.” Similarly, when opportunities have an expiration date, make that explicit so that teams and executives can have useful discussions about when it’s too late to go after an opportunity.

For many opportunities, risk is also a consideration. Another word for risk is *uncertainty*; uncertainty about the value of delivering a particular outcome, and uncertainty about the feasibility of delivering a particular outcome. A great advantage of delivering iteratively is that each delivery produces new information that reduces uncertainty and risk.

4. Continuously (Re)Evaluate and (Re)Order Opportunities

Different people in the organization will discover new opportunities all the time, and the relative attractiveness of existing opportunities will also rise, and fall based on new information. As a result, you will need to continually refine the relative order of your opportunities as you discover new information. This is the reason why you want to delay investing all your money as long as possible: you might find another opportunity that is even better.

To enable your organization to make the most of its investments, you will want to refine your opportunities so they represent the smallest amount of work that can be done to advance toward the full realization of the opportunity, typically the smallest possible experiment that a team can conduct to validate an opportunity. You’ll want to keep the opportunities up to date, so that a team could *pull* the opportunity to work on it at any time.

The order of the outcomes is important but arriving at a consensus is often a complex negotiation and consensus-building political process in which trade-offs between different competing business interests must be balanced. The best advice we can give to solving the prioritization problem is to be open and transparent about organizational goals and trade-offs, and work toward a consensus about the priority decisions.

The relative importance of pursuing an opportunity should be re-evaluated every time the organization has new evidence.

5. Minimize Avoidable Loss

To maximize the impact of your investments, you need to minimize avoidable loss by quickly learning which ideas *won’t* work, so that you can focus on which ideas will. You should think about funding product or service development in a series of experiments, each of which has the *goal* of trying to

prove that the solution won't work, so that if they succeed, they will improve your confidence that you're on the path to the right solution.

As an example, if you are worried that customers won't find your solution compelling, you could run a focus group with target customers using a mock-up product to evaluate their reaction to it. Or you could build a simple one-feature product and release it (easier to do with digital products than physical ones).

The idea is that if you can't prove that the assumption or concept for the solution is bad, it might still be good, and you continue with your experiments. If your experiment fails, you've not invested much, and you can move on to test other ideas.

To do this as quickly and inexpensively as possible, look for the simplest, fastest, and least costly way to empirically evaluate the workability of a solution; keep the size of investment as small as possible. In practical terms, this usually means looking for the greatest threat to the viability of the solution, then running an experiment to determine if this threat is real or just imagined. Ask yourself, "which of our assumptions, if it turns out not to be true, is most likely to kill this solution?" Then figure out a way to try to make your assumption fail, which forces you to see if you can find a way around it.

6. Focus on Improving the Flow of Work by Limiting Work in Process

Organizations generally have more ideas than they have capacity to pursue. When the organization tries to work on more ideas than it has capacity, they create a lot of Work in Process (WIP). This causes waste and delay when teams split their focus across more than one initiative at a time:

- Each initiative takes longer to deliver because teams lose effectiveness when they switch contexts between different initiatives
- Delaying value realization or feedback on value delivered deprives the organization of information it needs to make course-correcting decisions
- Delayed feedback causes organizations to waste time and money on things that later prove to be valueless once they get feedback

Having teams pull the most valuable opportunity only when they are ready limits WIP. When they do so, they are signaling that they:

- Have the skills to complete the work, and
- Have the capacity to do the work

When an organization has highly skilled cross-functional teams, nearly any team will be able to work on any opportunity. When a team is ready to pull work but lacks the right skills to work on it, the transparency the pull model provides helps everyone see where they team may need to develop new

skills or augment their skills by adding new or team members or adapting the composition of the team.⁵

Organizations also benefit when teams are able to decide what they work on, and when, because the freedom to choose helps them to be more committed. When teams pull their work, their performance improves because they focus on one thing at a time. In addition, they tend to be more motivated because they have an enhanced sense of autonomy.

7. Improve Status Reporting with Increased Engagement and Transparency

Traditional portfolio investments are monitored through a status reporting process that lacks transparency and consumes a lot of time and effort for little gain. When people outside the team are doing the work prepare status reports, you get a filtered view of reality. Subtle pressure to show that things are going well may mask important information that you need to know to make better decisions. Percent complete measures and Green/Yellow/Red status indicators don't tell the full story, and they may not even tell a true story of how things are going. The resulting lack of transparency leads to unpleasant but avoidable surprises if full transparency had been present earlier.

In place of filtered and subjective status reports, the results gathered from frequent, iterative product deliveries will provide you with real data - updated estimates of unrealized value based on evidence, as well as measures of current value delivered by the most recent release. Both will help you decide whether your assumptions are intact or whether your priorities should change. They will give you the information you need to know to reassess priorities and potentially realign funding with business objectives. As a result, you can improve your visibility into what is really going on by eliminating filtered status reports and replacing them with real evidence about how the initiative is achieving its objectives.

How to Get Started

Fully embracing all these principles, all at once, will challenge many organizations beyond their capacity to change; changing the way that investments are made and managed is a strategic function that forces many changes in the way an organization both thinks and works. To get started, treat the investment process itself just like one of the many possible investments in improvement that you can make: make small improvements, framed as experiments, make the change, measure the results, and inspect and adapt.

A good place to start is to sever the link between annual budgeting and project/product investments. This gives you the flexibility to decide how much team capacity you need and then later decide what you want those teams do using the Opportunity-driven process we've outlined above.

Converting all teams to an Opportunity-pull model that we've described is probably too much to do all at once. Start with a smaller Opportunity Backlog and a few teams and see how it works. Let teams get used to it, while you discover where your unique challenges lay. Add more teams and more opportunities as you gain experience and confidence. And learn empirically based on your experience.

⁵ For more information about how teams can improve their flow, see <https://www.scrum.org/resources/kanban-guide-scrum-teams>

There are no *best practices* you can copy; you are going to have to learn for yourself what works for your organization.

Conclusion

Opportunities, rather than projects or programs, provide a better way to align work with business objectives by focusing on delivering value to customers. Identifying and ordering opportunities based on estimates of *unrealized value*, letting teams *pull* opportunities to work on when they have capacity, releasing early and often, and measuring *current value* every time a new release is delivered gives organizations the flexibility and information, they need to make better investment decisions. Frequently re-evaluating funding based on best available evidence gives organizations a better way to balance competing demands, to more effectively manage risk, and to redirect their teams to go after better opportunities when they present themselves.