



The DIF: Designing For Lasting Impact That Goes To Scale

Mulago Foundation 2016

The DIF – Design Iteration Format – is a tool to help you design a model to create lasting change at a big scale and an organization that can take it there. The DIF process guides you through an initial (or remedial) iteration of design that can serve as the vehicle for a continuous process of refinement and evolution. It's really a sequence of steps, like this:

The Mission: exactly *what* you're setting out to accomplish.

Big Idea: *how* you're going to accomplish it.

Impact: the key outcome that fulfills the mission and focuses design.

Behavior: *who* must do *what* differently to create that impact.

Interventions: what you're going to do to drive those behaviors in a lasting way.

Model: a replicable, scalable process that integrates the intervention.

Doer & Payer: who will replicate and who will pay for your model at scale.

Scalability Audit: a simple and systematic assessment of the model.

The best way to use the DIF is to:

Read through the whole thing quickly to get a sense of the process.

Think long and hard about your mission. Everything downstream depends on it.

Do the same with impact. Get it right. It's the focus of the whole process.

Read each section carefully as you lay out each step.

Use the scalability audit to drive your first iterative tweak of the model.

Relax. You're probably going to change most of it anyway.

We'll illustrate the DIF process with the fictional work of a guy we'll call Jack. Follow him as he staggers through the process.

Jack throws nets out of a plane

Imagine this: our flight delayed, we meet a guy named Jack in the airport bar in Kampala. He's a garrulous guy in his 60's, made a bunch of money in his 40's and wants to "give back." Jack was a pilot in the Vietnam War and still flies a lot. He caught the Africa bug on a safari trip, and his adored grandkids have left him obsessed with the wellbeing of children.

Jack says he's started a project in a remote part of southwest Uganda. We brace ourselves. We think: *here we go again*. "I'm throwing mosquito nets out of a plane!" We sigh. Jack can see that we think he's nuts, so he starts telling us about the Poverty Action Lab study demonstrating that free distribution of nets is the best way to get kids under them, and how his analysis of the area showed him that lots of people are out of reach of government services or even roads. We perk up. Just then the loudspeaker announces that we'll be in the bar for four more hours, so we persuade Jack to be a guinea pig and go through the DIF process with us.

1. Mission

Decide on and write down exactly what you're trying to accomplish with your big idea. Everything you design and do will be shaped by it. Capture it in **eight words or less**, and include a **verb**, a *specific target* population or setting, and a big **outcome** that implies something to measure. The mission is about **what** you're going to do, **not how** you're going to do it. It for design, not your website. A good mission statement is generic – even kind of boring – but specific (and unvarnished), like this:

- "Get African one-acre farmers out of extreme poverty" (note: "extreme poverty" means not enough to eat.)
- "Prevent HIV infection in Brazil"
- "Save coral reefs in the South Pacific"
- "Educate slum kids in West Africa"

Be absolutely clear/explicit/specific about your target and your outcome! That is critical for a useful mission statement. Ambitious is OK, vague is not.

Jack's mission

Jack knows that children are the ones most at risk of dying from malaria, and that anything that benefits children will also benefit adults. So, he formulates his mission as:

"Prevent childhood malaria in remote regions of Uganda"

2. Big Idea

This is your central, distinctive idea about **how** to accomplish the mission – the idea at the core of all you do (some people might call this the center of your “theory of change”). Usually it’s about how you drive behavior. Get your idea down to a sentence that captures your special sauce. Keep working on it until you really like it. Capture your idea first if it helps you sort out your mission. This is a distillation of your idea to its essence in a way that will help you communicate it and think about it. Some examples:

- “Design, market and sell money-making products that farmers can afford and will use”
- “Eradicate devastating invasive species from islands so that endemic species and ecosystems can recover”
- “Use existing community groups to provide poor farmers with the integrated delivery of farm education, credit and access to cash buyers they need to make a decent living”

Jack’s Big Idea

Get kids under mosquito nets by delivering insecticide-impregnated bed nets in high-malaria regions that can’t be reached by road.

3. Impact

Impact = change in an outcome. Identify the *single best* outcome – an outcome, not an activity, not a behavior – that would let you know if you’re fulfilling the mission. This makes you focus in on a specific, granular and practical outcome. Yes, it’s hard to narrow it down to one, but do it anyway. It needs to be at least theoretically observable and measurable, even if perhaps really hard to do in the real world (Don’t worry about actual measurement yet). This is a critical step: *it defines and focuses what you’re designing for*.

Here are some examples from the missions above:

- “Increase in *farm income*”
- “Decrease in *HIV infection rates*”
- “Improved biological indicators of *coral health*”
- “Functional *literacy and numeracy*”

Jack’s impact

Decreased *malaria rates in children under five*

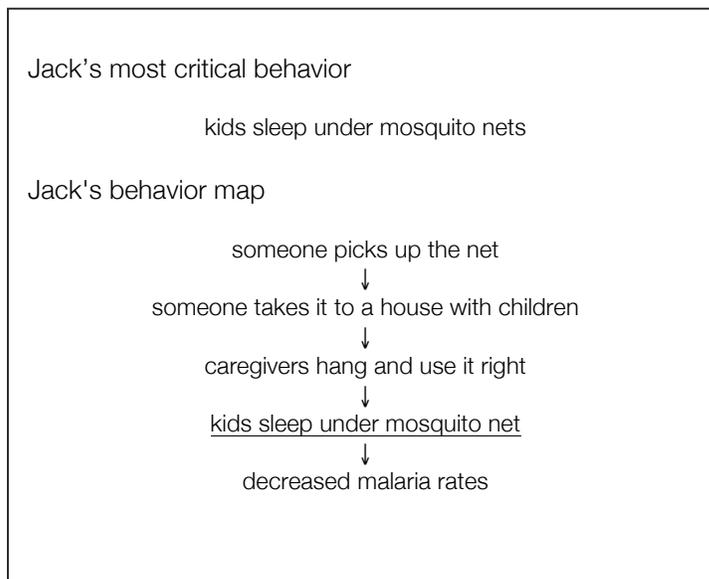
4. Behavior Mapping

Impact comes from action, from someone doing something differently. Given your idea, **who** must do **what** to create impact? Design for impact is focused on behavior – what must happen, how to drive it and how maintain it.

To start, put down the most critical behavior change that must happen for impact.

- “Farmers adopt new set of farming practices”
- “Teenagers practice safe sex”
- “Island communities guard reefs”
- “Teachers show up for work”

Behaviors that directly create impact don’t happen in isolation, though – getting to impact usually requires a connected sequence of behaviors. To get to impact, you’ve got to “connect the dots” and run through all the behaviors necessary to get there. List them in sequence as a simple flow diagram of *who must do what* all the way to impact. Start with outcome, add in your most critical behavior and construct the sequence.



A behavior map is about what people must **do**. It’s *not about attitudes, thoughts, or awareness* – it’s about behavior. It’s only about those who you can’t fire; in other words it’s not about your employees or others you can directly manage. It’s about those whose behavior you and your organization have set out to change.

5. Interventions

Look at your list of behaviors. Figure out which ones would happen without you and take them off the list. Of those that remain, put down – in brief – what you’re going to do to

make them happen. These are your *interventions*. This is the list of the things you have to do – the activities that you must craft into a scalable model.

Bear in mind that changing behavior is hard, and that ideas that require you to change more than 3-5 behaviors are going to be pretty complicated (i.e., not very replicable).

This is the step that often sends people scurrying back to do more research as it becomes clear that they didn't think about some aspect of behavior change.

Here's something worth doing once you've listed your interventions: Look at each necessary behavior in terms of **conditions and incentives**: are the conditions in place such that the behavior *can* happen, and are the incentives there so that it *will* happen? This is a systematic way of checking that the intervention is sufficient and it often suggests ways of adding to or modifying it. Can happen, will happen, can happen, will happen: that's your mantra. A good starting point for **can happen** is "make it easier to do it than not do." A good starting point for **will happen** is "link it to something people already want."

Jack's interventions

Uh oh. Jack now realizes that he isn't certain that any of those behaviors will happen if all he does is throw nets from a plane. He makes a desperate stab at some initial ideas:

Someone picks up the net:

- Map and target well
- Radio announcements, posters at markets
- Bright packaging

Someone takes it to a house with children:

- Radio announcements, mobilize churches, posters
- Saturation – excess nets until market is saturated

Caregivers hang and use it right:

- Radio, posters, sermons
- Effective instructions in packaging
- Net design tweaks

Kids sleep under it:

- Radio, etc.

Those are just some initial ideas. Now Jack has to do some homework and he has to iterate: This is where user-centered design, rapid prototyping, A-B testing and all that stuff comes into play. Somebody sends Jack a good behavior change booklist to :

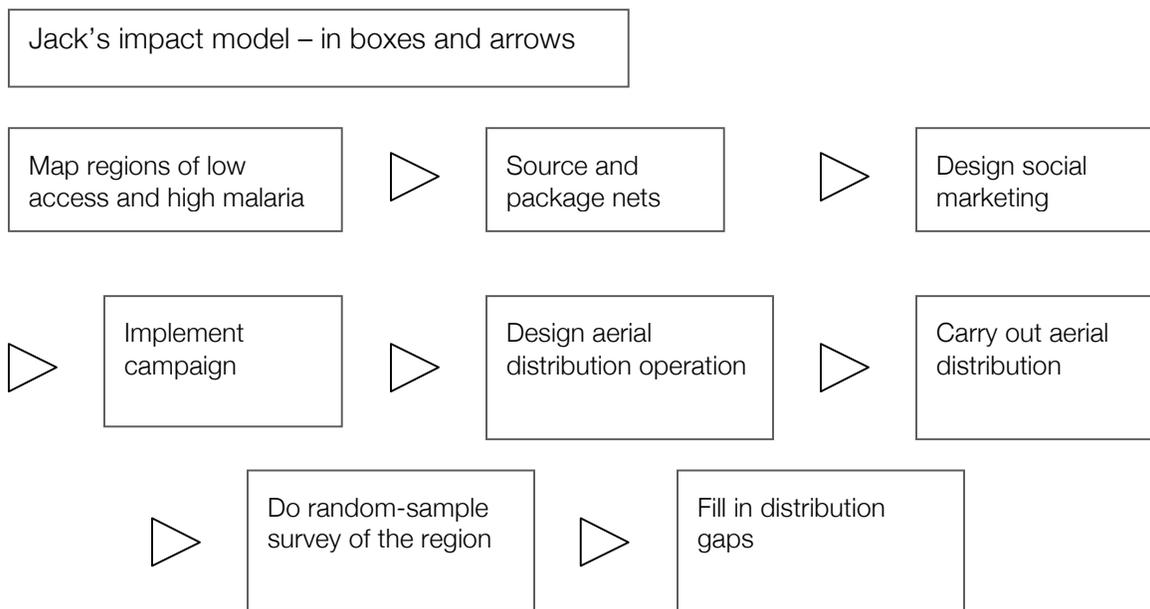
Swth. Heath and Heath
Nudge. Thaler and Sunstein
Thinking, Fast and Slow. Kahneman
Predictably Irrational. Dan Ariely
Influence. Cialdini

That turns out to be a good start, and Jack heads down the iterative rabbit hole of design for behavior change.

6. Model

Nothing is replicable unless it's systematic. Any doer – you or someone else trying to replicate what you do – needs a systematic operational model to follow. This is a way to get you thinking systematically inside a format that allows you to continually mess with your model. Don't worry about the exact form – just try it. You can use a boxes-and-arrow format, some other kind of flow diagram, or simply a list. *Use a pencil and blank paper – or post-it notes or whatever – if that works better than doing it on a computer.*

First, list all of your interventions. Then try to integrate them into a minimal number of discrete steps in an operational. There is no prescribed right way to map out the boxes and arrows; just keep fiddling with it until you have something that makes works for you. The result is a flow diagram of your impact model. Look back and forth between Jack's interventions above and his model below to get a sense of how it translates.



Another option is to fill it out in list form:

- Map regions of low access and high malaria →
- Source and package nets →
- Design social marketing campaign →
- Implement campaign →
- Design aerial distribution operation →
- Carry out aerial distribution →
- Do random-sample survey of the region →
- Fill in distribution gaps

However you do it, make it describe a systematic process of sequential steps that integrates all of your important interventions: If you or someone else was to replicate somewhere else, these are the steps they'd need to take.

The devil is, of course, in the details. Once you've sketched your impact model to capture a replicable process, you'll probably find it useful to articulate a brief narrative that captures what is distinctive and necessary for each step. Imagine that you're trying to describe what you do to someone interested in doing it for herself. Don't put too much effort into this; just capture the detail that makes it actionable.

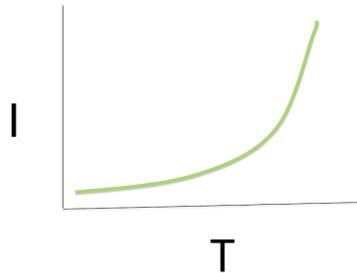
Some of Jack's detail

Design social marketing campaign

Systematically figure out what channels exist: radio, other media, churches, markets, other widespread groups, grass-roots communication, key social networks. Explore the possible role of each channel in each necessary behavior. Integrate promising ideas into a cost-effective campaign.

7. Doer & Payer: Who's going to take it to scale?

You've figured out your model, and you're serious about impact at scale. Scale is more than just growth. It's this:



Call it exponential, geometric, what have you – the point is that the curve steepens and impact accelerates dramatically over time. It's “getting to a million” – a million lives saved, kids educated, etc. – and on to subsequent orders of magnitude.

If you want to get to real scale, there are two critical questions that really matter: 1) who's the **doer** and 2) who's the **payer**? You've got a model – the thing you do. Who's going to replicate – do – it and who is going to fund that replication?

The Doer.

When you eventually get your model dialed in – proven, tuned, and through initial rounds of replication – somebody has to replicate it at scale. There are only four choices:

You: running an *NGO or business* that gets to scale through growth or leverage
Lots of NGOs: replicating your model
Lots of businesses: replicating your model
Governments: delivering your model through programs and policies

That's all you get. Pick the doer that will dominate at a scale of a million and beyond. They've all got pluses and minuses, like this:

You: Having full control over replication means that you can deliver a complex model at high quality. Building and growing a really big organization is a pain in the ass, especially in a dysfunctional funding market.

Lots of NGOs: Plenty of bandwidth there, and it shifts fundraising off your back, but NGOs are notoriously bad at implementing other NGOs ideas. NGOs that grow really big usually can't maintain disciplined replication of one solution because they have to meet a huge payroll.

Lots of businesses: One-off businesses are rarely enough – it's industries that solve problems. The market has nearly infinite capacity for replication and businesses can replicate fairly complicated stuff. Obviously, to make use of this doer, a solution needs to come with a profitable business – the more profitable, the more imitators – and that precludes a lot of solutions. And too often for-profit solutions leave out the very poor.

Government: They have big bandwidth, lots of resources, and a mandate to serve – and they're probably the only way a lot of basic service solutions will scale. They are usually inefficient, maddeningly inconstant, and often corrupt. They have a hard time delivering anything very complicated.

That's it. Pick one. A model will only scale if it is designed with the doer in mind. People do attempt to replicate via multiple doers at scale, but one usually dominates over time (Or the whole thing fails because multiple doers diffused your efforts). Start with your best guess and design toward it. Like everything, the process is iterative: you may change course if becomes clear over time that another doer is better.

Got it? OK, now pick a payer at scale. Somebody has to fuel replication. Conveniently, there are only four payers as well:

Customers: revenue from product sales

Taxes: revenue raised by governments

Big Aid: multi- or bilateral revenue from rich governments to poor governments; sometimes delivered straight to the doer

Private philanthropy: anyone from the big foundations to the individual small donor

Pick the one likely to be the *main* source of funding at a million and beyond. Again there are pluses and minuses:

Customers: Investment gets a business growing, but eventually you need to sell a lot of stuff to a lot of people. To fuel scale, the businesses have to be either high-margin or high volume. Low margin businesses only scale if they supply a widespread demand.

Taxes: Most poor governments don't collect much in the way of taxes. Much of what they do is ultimately funded by Big Aid. That's tricky to put together and coordinate over time. In poor countries Taxes and Big Aid often come as a package deal.

Big Aid: Big Aid is big, but it's a huge hassle and usually comes a lot of strings attached. Priorities change often.

Private philanthropy: Relatively small. Quirky. Can only drive scale if the model is high-leverage – digital/mobile solutions sometimes provide enough efficiency/leverage. Usually better at earlier stages of growth.

Gather your courage and try to fix on one doer and one payer. That will give you the best chance at scalability.

Jack's doer: government Jack's payer: Taxes/Big Aid

By the way, here is Jack's thinking: "We will build an organization big enough to conclusively prove the model's impact and cover a substantial part of Uganda, then we'll leverage the results into policy change and an effective government program. Once we've got effective take-up by government in Uganda, we'll begin to lobby other governments to establish similar programs and provide the technical assistance they need." (NB: you don't need to write this out here – just which of the four you're going to use.)

8. Scalability Audit

Once you've got your doer/payer figured out, there are three big questions that help you assess the fundamental scalability of your model:

1. Does it apply enough?
2. Is it simple enough?
3. Is it cheap enough?

Apply enough: Are there enough settings where your model would apply, where it would plausibly work? There is going to be a minimum set of conditions that must obtain for any model: Are there enough settings with those conditions? Is it adaptable enough to work well in those settings? Cookie cutter solutions don't work very often. The best models use a systematic process to generate locally appropriate solutions. If your model is rigid, or if it depends on a limited set of circumstances or a rare kind of talent, then it's not going to get very far beyond local scale.

Simple enough: Could your doer do it at high quality? Could they plausibly replicate well enough to get the same impact? If your model is complex and requires ninja-level execution, then you're going to be doing it yourself. Governments are pretty hopeless at it, and NGOs are often not much better. Even businesses won't be able to replicate a model that is very complicated.

Cheap enough: Is it cheap enough for your payer to pay it? Everybody has a price point. This is about cost-effectiveness – the cost per unit impact – but it's also about the price that the payer is willing to pay. If you want to scale up your innovative community health worker model, what's the Health Ministry's price point for per capita coverage? If you have a product to sell, will customers pay enough to give you (and your imitators) a decent margin?

To be scalable the answer must be yes to all three. A 30,000' assessment is useful, but getting to real answers might require some serious homework and projections. Like everything else here, the answer to these questions is iterative. Revisit them often and continually tune toward a more emphatic *yes!*

Finally.

Iterate, Do, Repeat.

Go.