The Nudge Unit

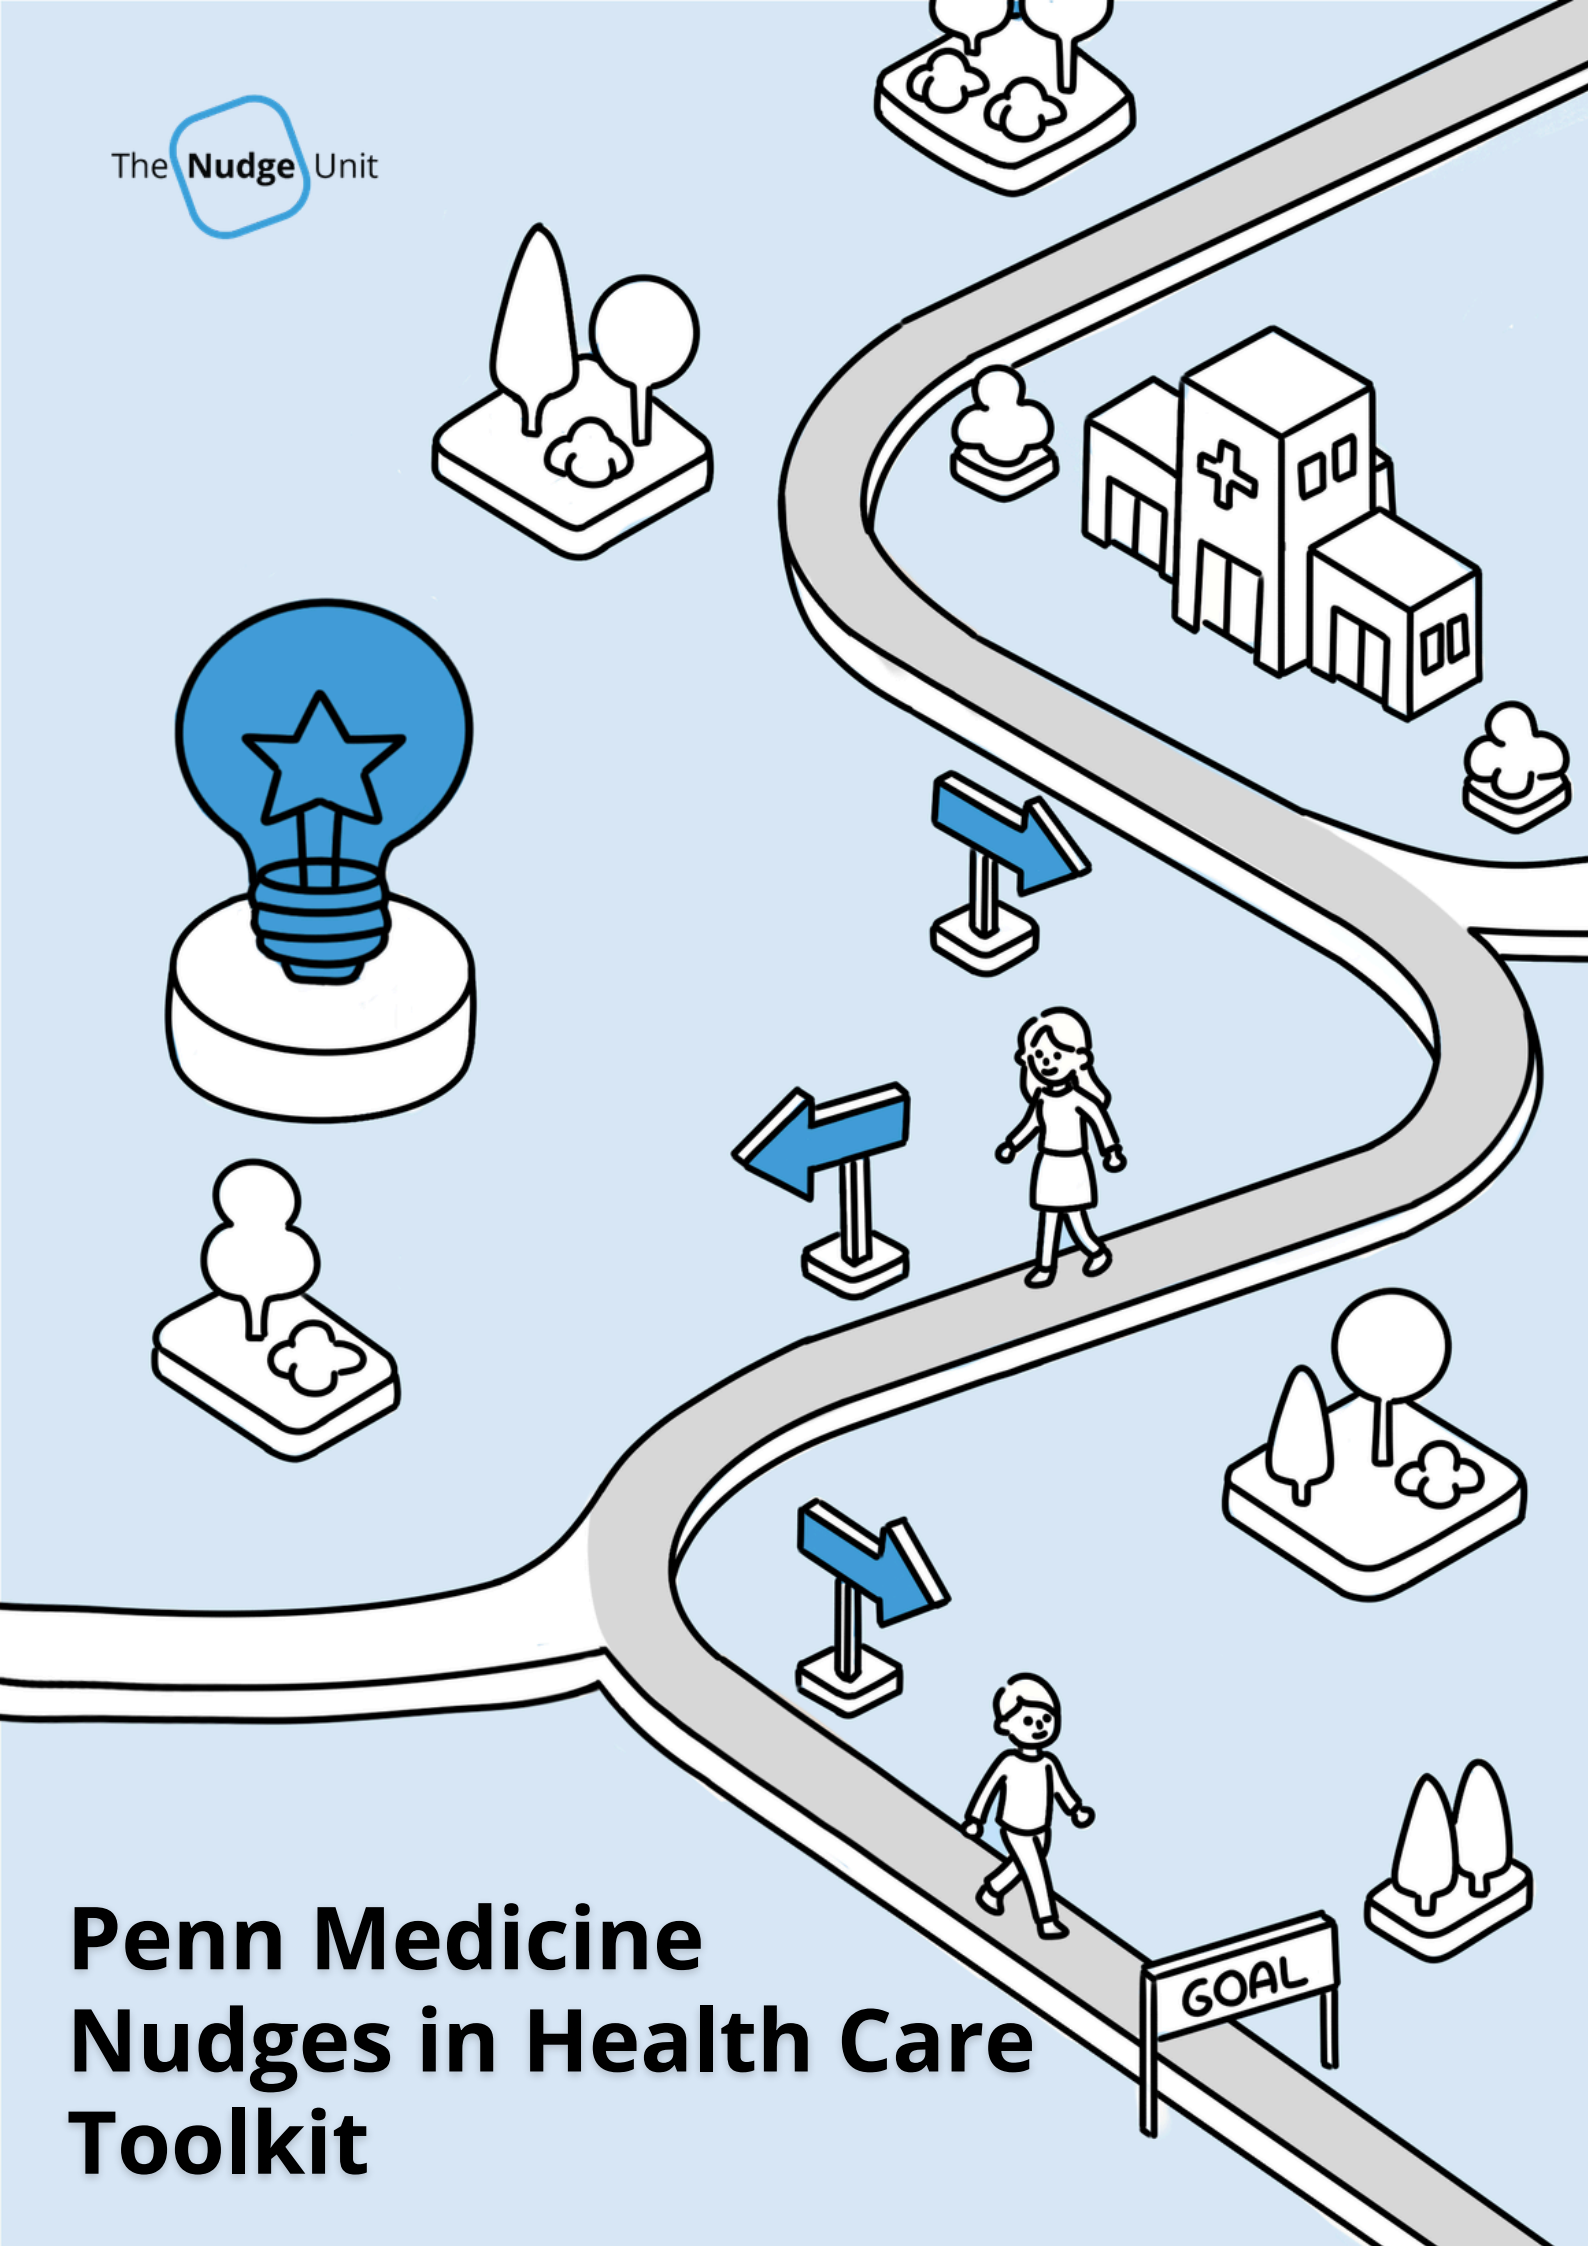

Penn Medicine Nudges in Health Care Toolkit

Table of Contents

Introduction	03
Kickoff	07
Discovery Research	16
Understanding and Changing Behavior	22
Designing for Scale	29
Prototyping	33
Evaluation Plan	37
Build and Scale	46
Conclusion	54

1
Introduction

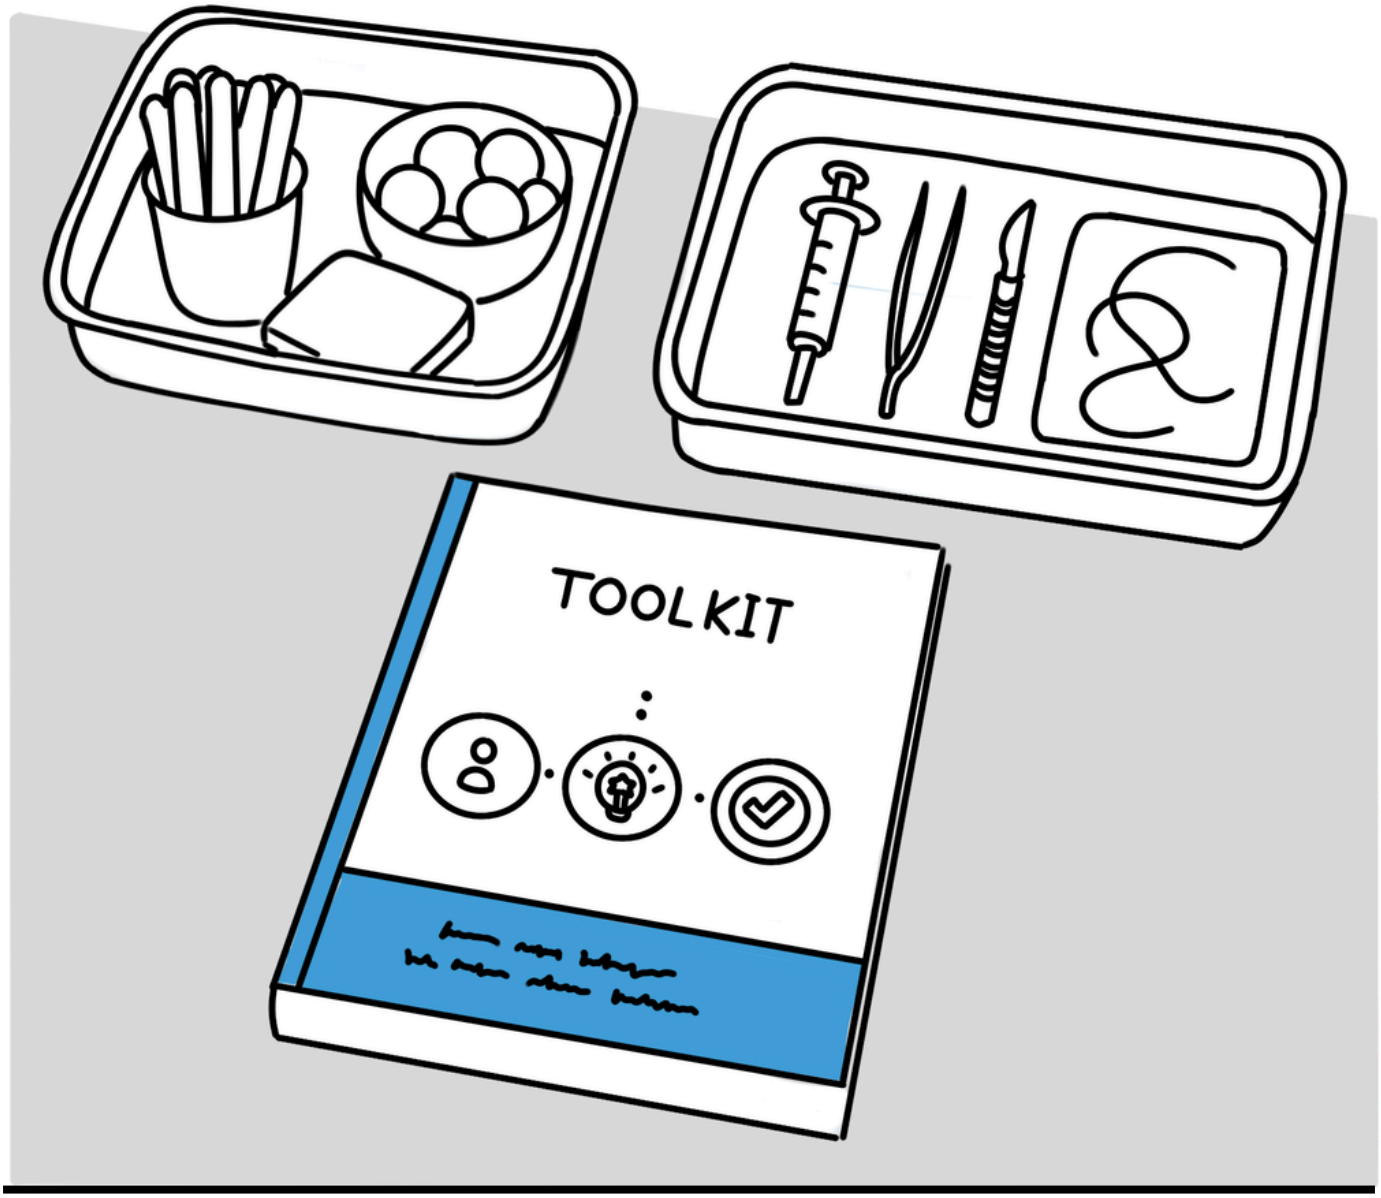

The Penn Medicine Nudge Unit

The Penn Medicine Nudge Unit is the world's first behavioral design team embedded within a health system.

The Nudge Unit is supported by the Penn Medicine Center for Health Care Transformation and Innovation and the Penn Center for Health Incentives and Behavioral Economics. *https://chti.upenn.edu/nudge-unit*

Our mission is to design, implement, evaluate, and scale evidence-based, behaviorally informed nudges that steer decisions toward higher-value care, better patient outcomes, enhanced public health, and greater health equity.

We work toward a future where the right choice is the easy choice for health.

What Is a Nudge?

A nudge guides or motivates decision-making by changing how information is presented or choices are framed. A well-designed nudge will make the right choice the easy choice and help people act in their own best interests. Whereas laws, rules, or mandates seek to force behavior, nudges always preserve freedom of choice.

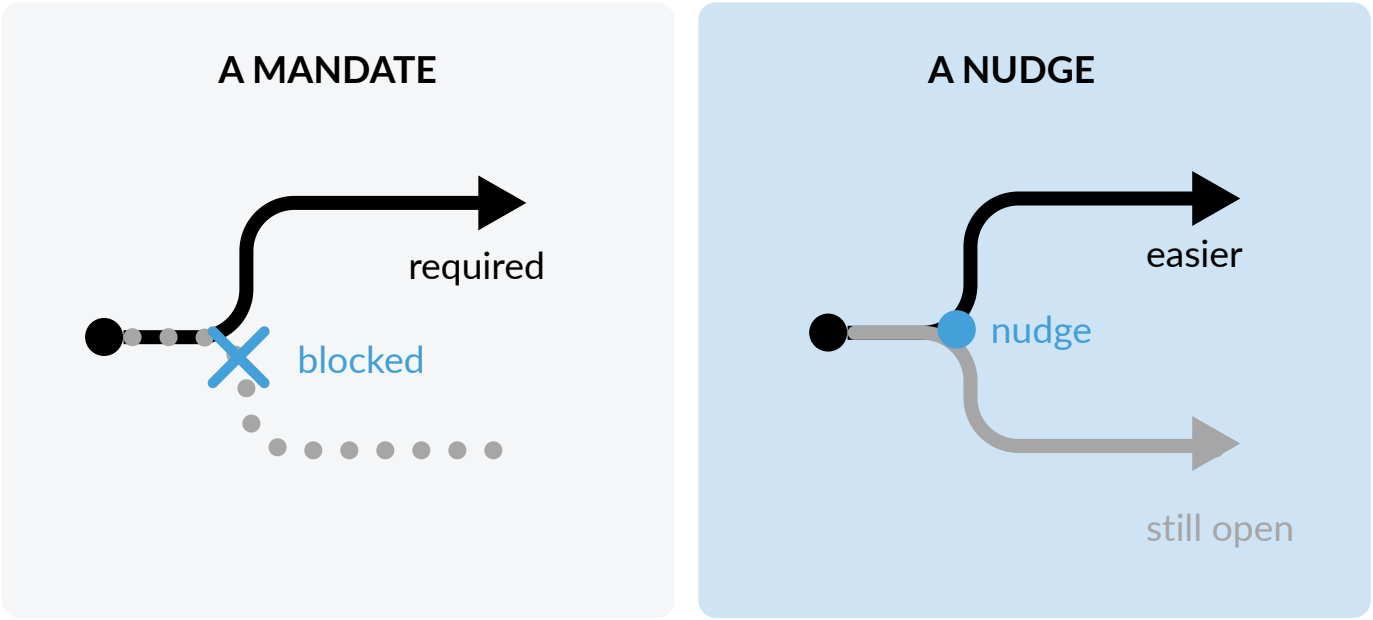

Which Problems are "Nudgeable"?

Nudges aren't the right choice for all problems. A good candidate for nudging has many of the following characteristics:

✔ A GOOD CANDIDATE FOR A NUDGE

- There is a **measurable gap** between where things are and where they should be.
- There is an **established best practice.**
- The best choice is **unclear to users** in the current situation.
- The current **environment makes it difficult** to make the best choice.
- There is an opportunity to **boost motivation** among care teams or patients.
- **There aren't currently any prompts** to trigger a desired action.
- There is **buy-in from key decision-makers** to enact change.

✕ PROBLEMS WILL NEED MORE THAN A NUDGE

- When there is **no established best practice to "push"** users towards.
- There is **no buy-in from key decision-makers** to enact change.
- **Extensive structural redesign** is required to achieve outcomes.

Why Nudge in Health Care?

Human behavior is a vital piece of the puzzle when it comes to advancing patient care. No matter how effective a medication, protective a vaccine, or beneficial a lifestyle modification is, they can only improve health if clinicians and patients understand and utilize them.

This is where nudges come in. A nudge is a change in the way information is presented, or choices are framed, that is designed to prompt, guide, or motivate good decision-making. The idea is to make it easier to do things you would want to do anyway. In the end, this reduces burden on busy clinical teams and helps guide patients through the confusing health care system.

How to Use This Toolkit

The Penn Medicine Nudges in Health Care Toolkit is designed to help a wide range of individuals and organizations implement quality improvement projects. This guide will walk you through the Penn Medicine Nudge Unit's process so you can design evidence-based nudges in your own work.

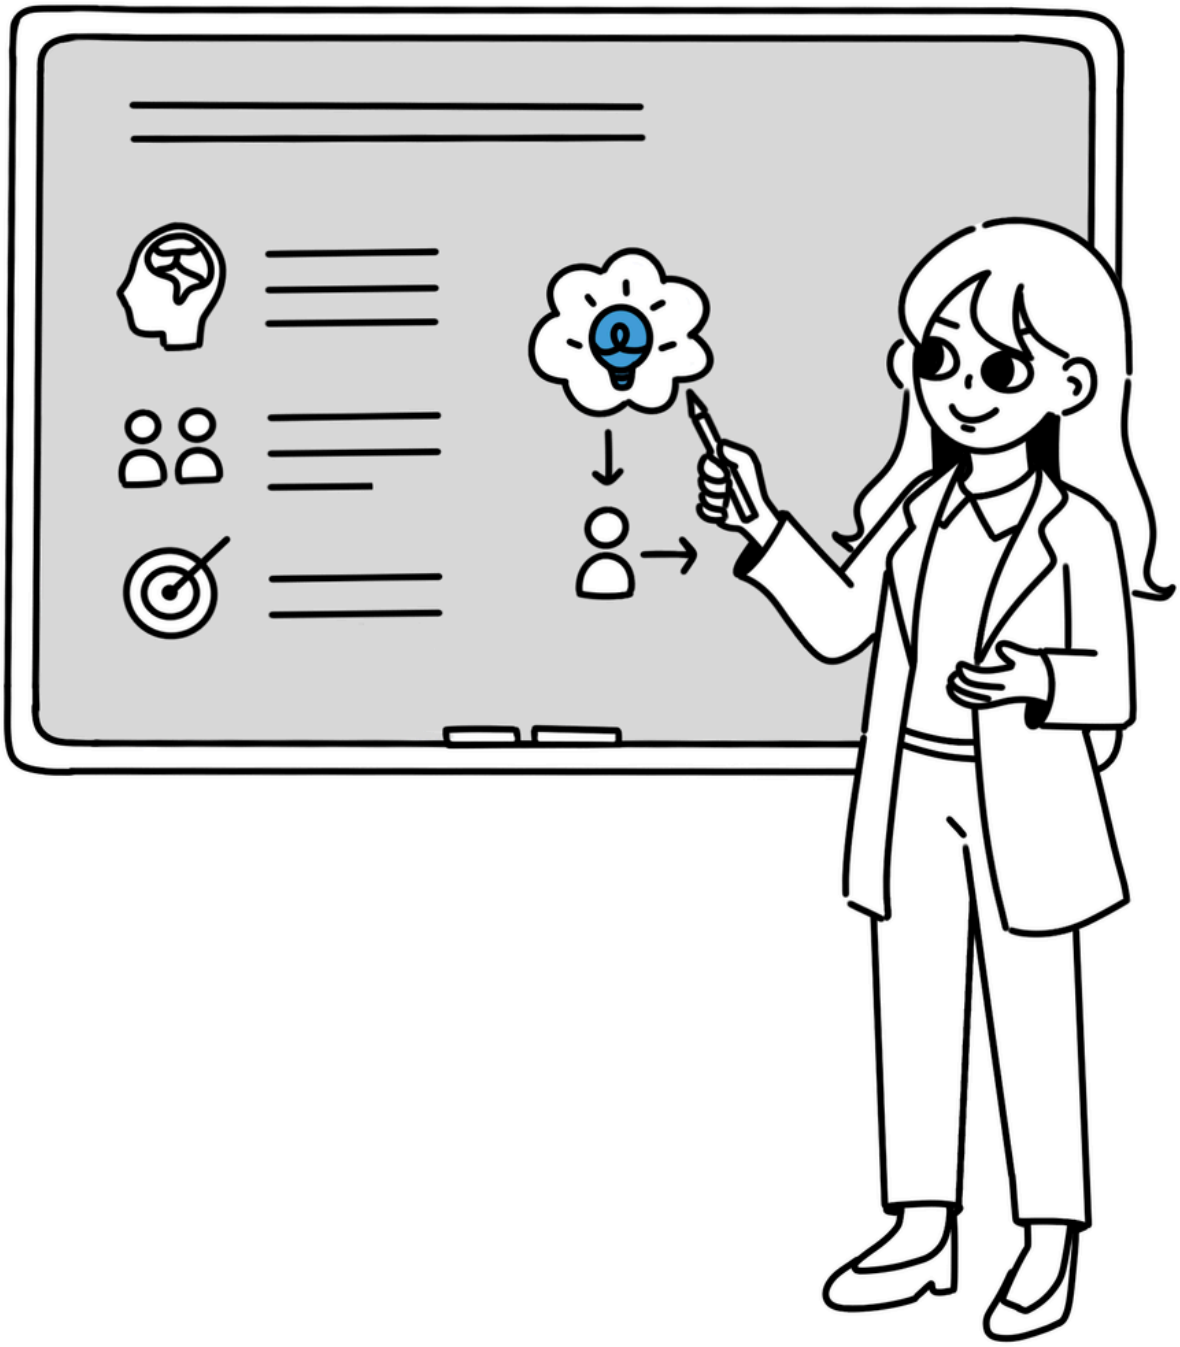

2
Kickoff

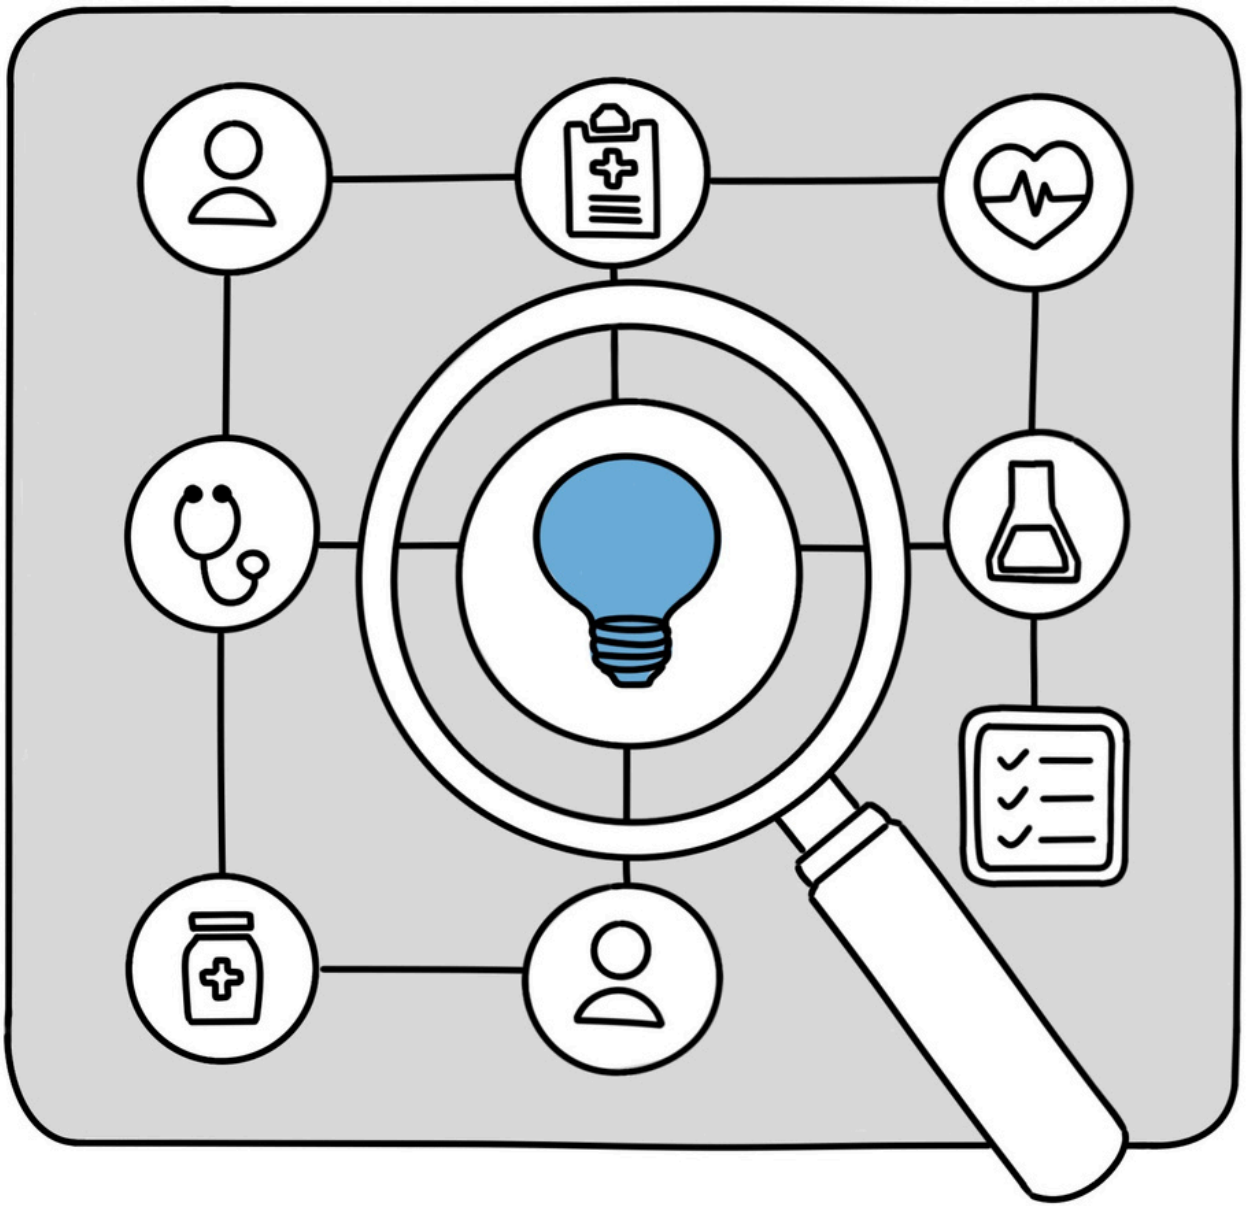

Defining the Problem and Target Audience

In the kickoff phase, it's important to clearly understand the problem you're trying to solve. A problem statement defines and frames that problem and sets the scope for your research. It describes what you're exploring and what you're aiming to achieve.

When you skip defining the problem and target audience, everything downstream suffers:

Solving the wrong problem & stakeholder misalignment

Starting work without clear direction leads to meandering research, wasted effort and resources, and declining team and stakeholder morale.

Building the wrong solution

When you start with a solution rather than the underlying problem, you risk building something that doesn't address the real need.

No way to measure success

Without a defined problem, there's no baseline or target to evaluate whether your solution accomplished the goal.

Unfocused research

Without a target audience, insights become too broad and difficult to act on.

Gather your team and work through the following questions together to begin shaping your problem statement.

The 5Ws

To gather relevant facts for your problem statement, you can use a simple technique called the 5Ws, which involves answering the questions below.

Next, to build your problem statement, combine these answers into one or two sentences. Here's a basic formula.

WHO	Who is affected by the problem?
WHAT	What is the problem?
WHERE	Where does this problem occur?
WHEN	When does the problem occur?
WHY	Why does the problem occur?

[Who] is experiencing **[What],** at **[Where]** during **[When],** resulting in **[Why].**

Example: Urine Drug Screening
*[**Clinicians**] are **[uncertain about urine drug testing guidelines]** on **[the labor floor] [when screening patients]**, resulting in **[unnecessary and inappropriate drug tests].***

Your problem statement will likely change.

As you gather more information, stay open to refining your problem statement — it will likely change and that's expected. Keep your problem statement front and center as you develop your nudge. Write it out on an index card at your desk. Pin it to your team's messaging channel. It helps tell you where to focus next.

Jobs-to-Be-Done Canvas

It’s likely you bring expertise in a certain area to this problem space — like a clinical specialty, health care operations, or other experience. This makes it easy to feel as though you know the problem best. However, you are not designing for yourself. Solutions designed for oneself are likely to fail.

> **Your users are the experts, the source of truth, and the starting place for designing any solution.**

Even if you yourself will be a future user of the nudge, your idea will need to be adopted by others to be successful.

That's our next task: exploring who is directly affected by the problem. There could be multiple users impacted by the problem in different ways. It's important to identify who they are.

Who are we designing for?
What are their goals and values?
What gets in their way?
What helps them succeed?

Jobs-to-be-Done is a framework that groups people by shared problems and needs — called "jobs" — which can be functional, emotional, or social. This is commonly used when designing digital products.

WHY NOT START WITH DEMOGRAPHICS?

If you often start a problem with demographic data, this framework may feel new. Demographic data is useful—it can make sure a sample matches the population or show differences between groups. But these labels don't decide how someone acts. People of different ages, races, sexes, and gender often share goals, motivations, and fears. We want to find those shared needs and put them first in our design to make the biggest difference.

Jobs-to-be-Done Questions

Customer jobs	• What goals is your user trying to accomplish in work, or life? • What problems are they trying to solve? • What needs are they trying to satisfy?
Pains	• What gets in their way? • What keeps them up at night? • Includes risks, bad outcomes, annoyances
Gains	• What benefits does your user want? • Includes functional benefits (ex. health outcomes), but also social benefits, positive emotions.

Process Mapping

Next, we'll look at the "when" and "where" answers to your 5Ws by process mapping. This helps us visually map workflows. By outlining a process from start to finish, we'll gain a better understanding of how the entire workflow operates. It gives you more precision in creating your nudge. According to behavioral science literature, nudging the right person, at the right time, is critical to a nudge's effectiveness.

Process map instructions

1. List out all the activities involved in sequential order.
2. Create a horizontal map of these activities, using the process map symbols.
3. Identify pain points and opportunities.
4. Specify when and where the problem occurs.
5. If you don't have an idea for a nudge already, look at the opportunity areas for the best place to put a nudge.
6. Create a future state, proposed workflow with your nudge idea included.

Process mapping symbols

Activity:
Represents a step, task or activity performed by a person

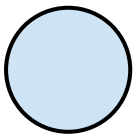

Start/End:
Represents a trigger that starts/ends a process

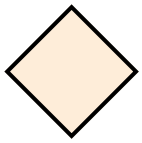

Decision:
Represents a decision point, where the flow will branch into two or more paths

Flow:
Indicates the order of events

Keys to use on the workflow

Pain points

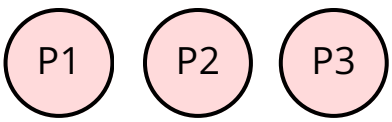

Opportunity for nudge

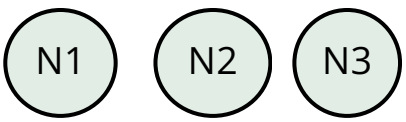

Process map example

Urine Drug Screen (UDS) current workflow

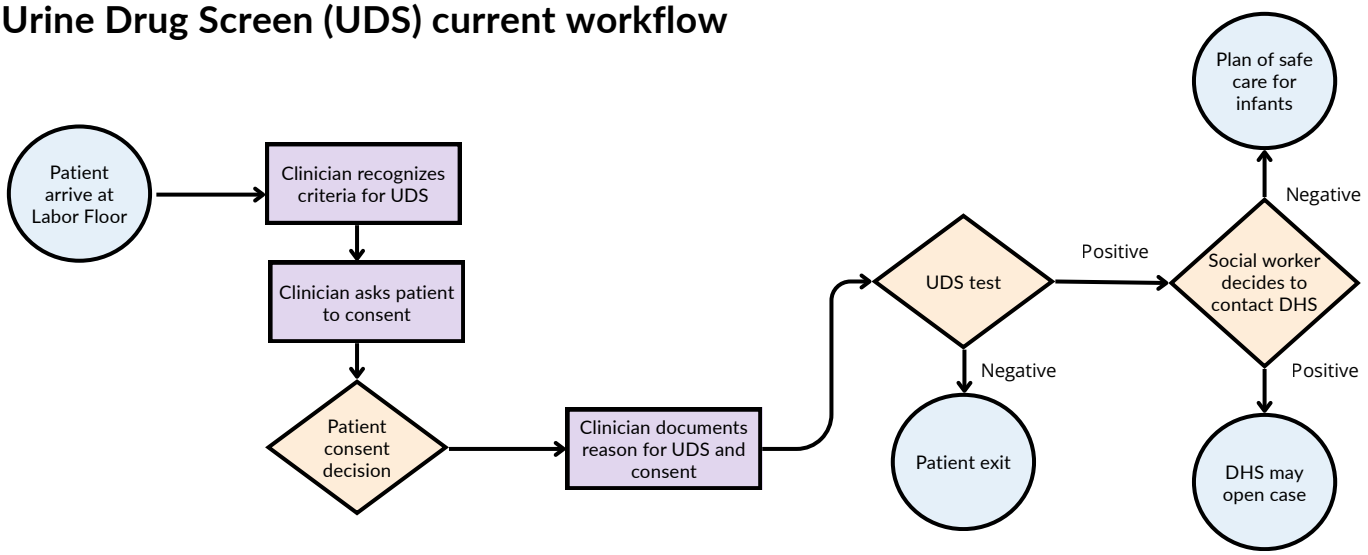

Urine Drug Screen (UDS) proposed workflow

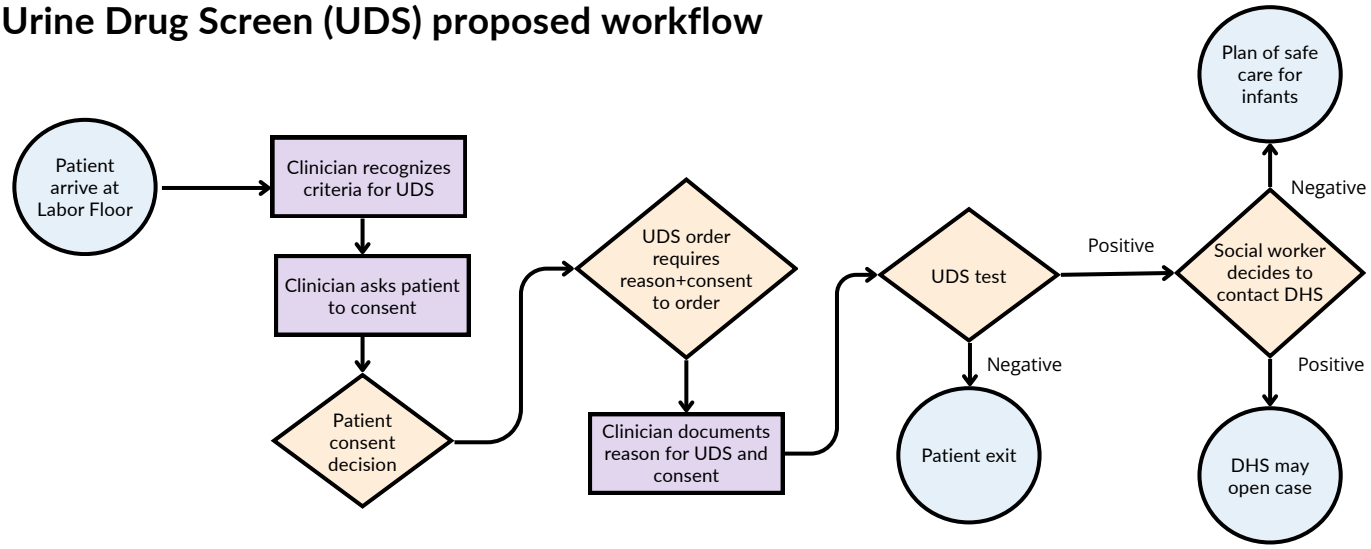

Assumption Mapping

Assumption mapping of proposed workflow

Now that you have your process maps, we can de-risk your solution. Assumption mapping helps you look at your solution with a critical lens. In this exercise, an assumption is anything that must be true for your solution to work. By identifying assumptions on your proposed workflow, you will surface potential issues early, allowing you to test them before investing in a solution. This prevents costly, time-consuming effort redoing work later down the line.

Doing this activity as a group brainstorm will generate more assumptions, or ones you might have missed. If you are working solo, ask people who have knowledge of the problem space to join you.

The process starts with a key question:
What must be true for this solution to work?

Example assumptions:

- *"Residents feel comfortable consenting patients."*
- *"Patients understand the reason for the order and opportunity to opt-out."*
- *"Clinicians know how to access in-person or virtual interpreters."*
- *"If a clinician sees this alert, they will act on it and not close it before reading."*

Instructions

1. Grab sticky notes, a large-scale version of the process map, Sharpies, a timer, and your fellow brainstormers.
2. Review the proposed process map steps, and suggested nudge as a group.
3. Brainstorm individually for 5 minutes.
    a. Write the assumption as a full sentence, with one assumption per sticky note.
    b. Add to the assumptions sticky note directly to the workflow at the time in the workflow where they would be relevant.
4. Discuss assumptions and cluster similar themes.

Adding assumptions to the workflow

Assumptions shown in red.

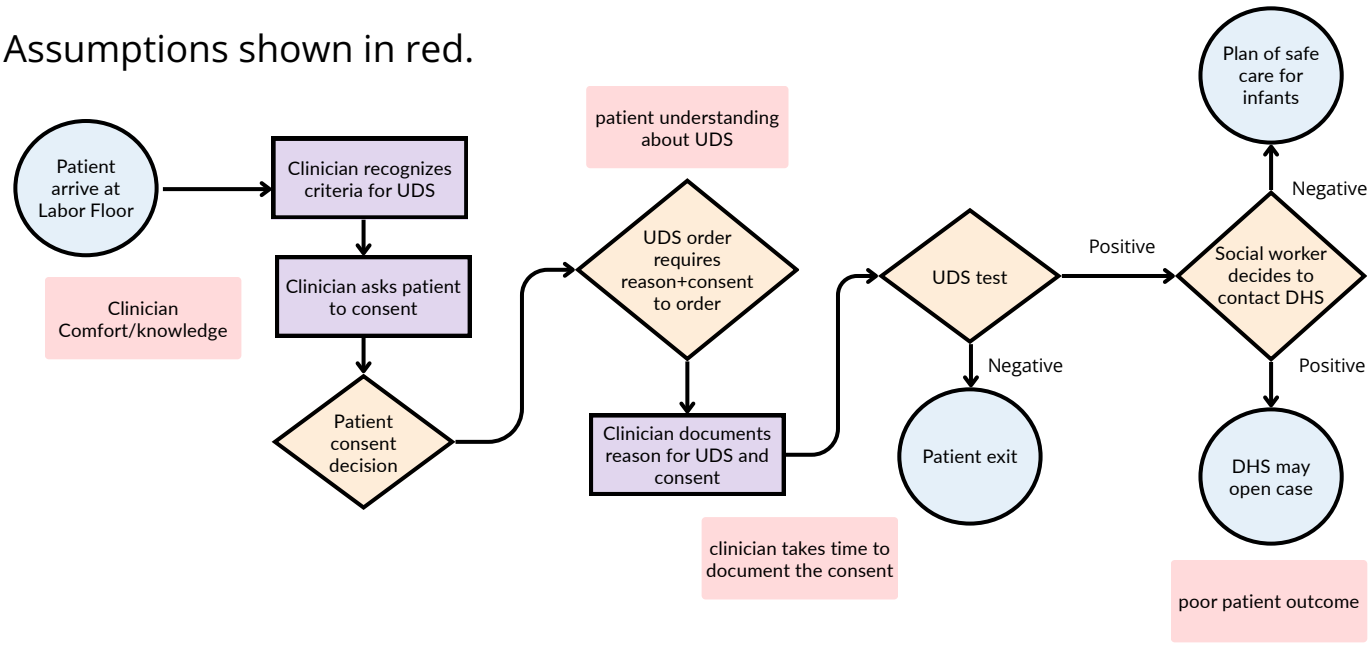

Risky Assumption

Once all assumptions are surfaced, plot them on a 2x2 matrix with two axes: how certain you are that an assumption is true, and how much is at stake if it isn't. This creates a clear picture of where your risks live.

The most critical area is the top-right quadrant — assumptions that are both high-stakes and low on supporting evidence. These are your **risky assumptions**, and they deserve the most attention. From there, identify your **top three riskiest assumptions** and prioritize validating them first, as they have the greatest potential to make or break your nudge if left unaddressed.

Instructions:

1. Review your assumptions for high stakes, low certainty assumptions.
2. Collectively place the risky assumptions in the 2x2 matrix.
3. Discuss the top riskiest assumptions.

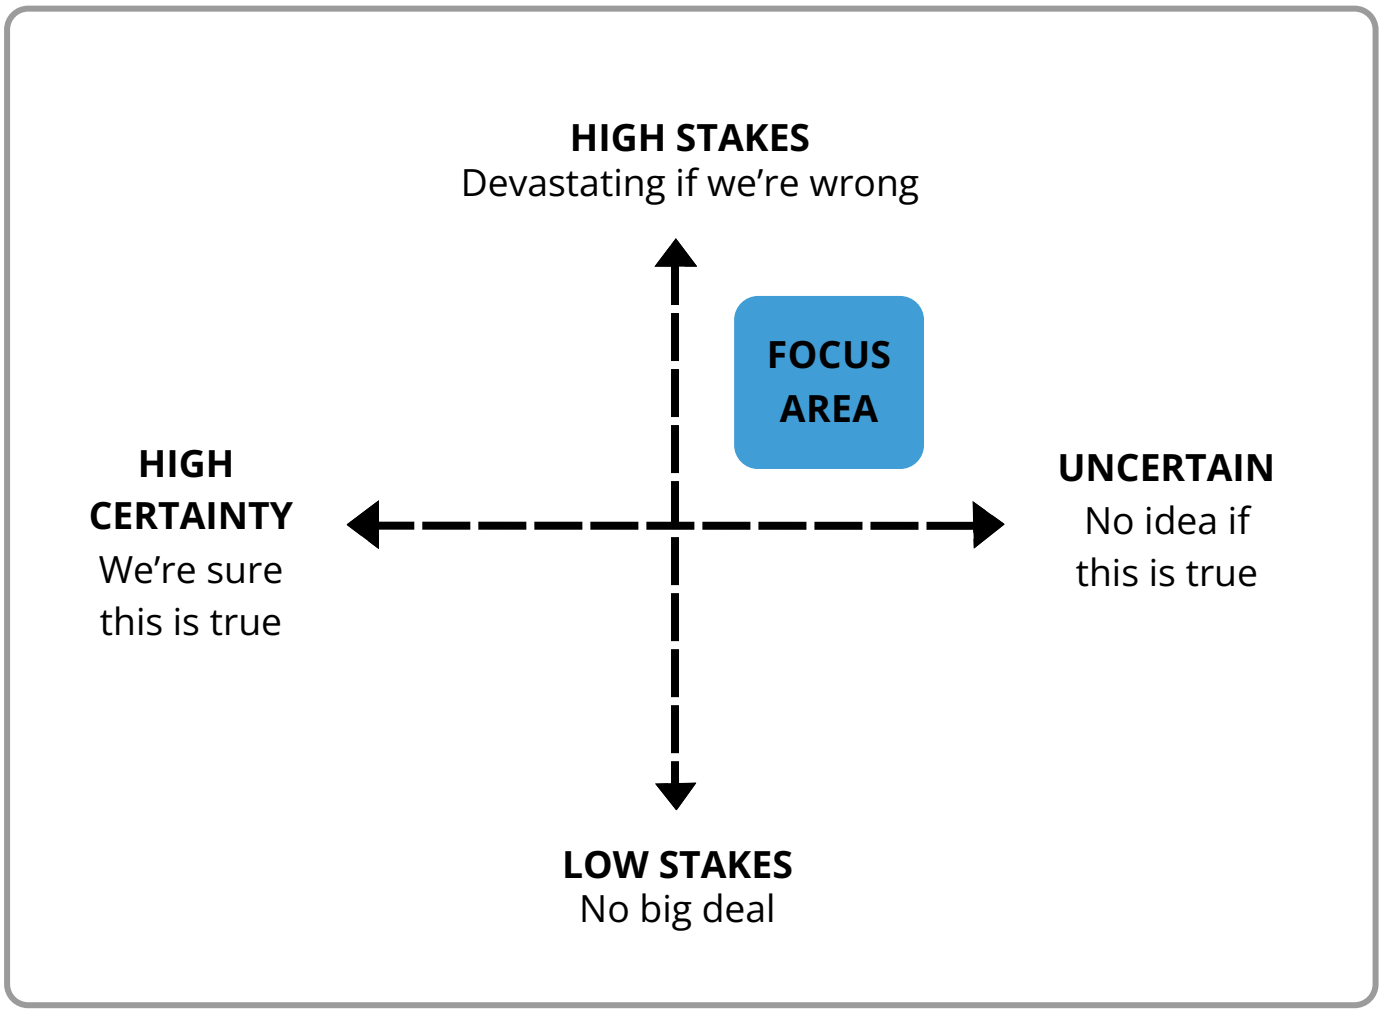

3
Discovery Research

Now that you've identified risky assumptions, the learning goals for discovery research are clear. During this phase, pay extra attention to who and when in the workflow you aim to nudge. This is critical for designing an effective nudge. If you nudge the wrong person, or nudge at the wrong time, even the best nudge will fall short.

Contextual Inquiry

Contextual inquiry, a combination of interviews and observations, is an excellent way to investigate risky assumptions. It involves focusing on a small sample of users to gain a rich understanding of their attitudes and behavior. Researchers watch users perform real tasks and ask questions along the way to understand not just ***what*** people do, but ***why***. This makes it easier to spot habits and details that users themselves might not think to mention.

In general, recruit five participants per role. In a health care setting, it's likely any workflow affects multiple roles so you will want to seek out multiple people. If you go above 15 people, you will often notice that fewer new themes emerge. There's no added benefit to keep interviewing beyond that point. In most settings, you should reach your target sample in 1-2 visits. With hard-to-reach users or very busy clinics, you may need additional trips.

NEW THEMES DISCOVERED PER ADDITIONAL PARTICIPANT

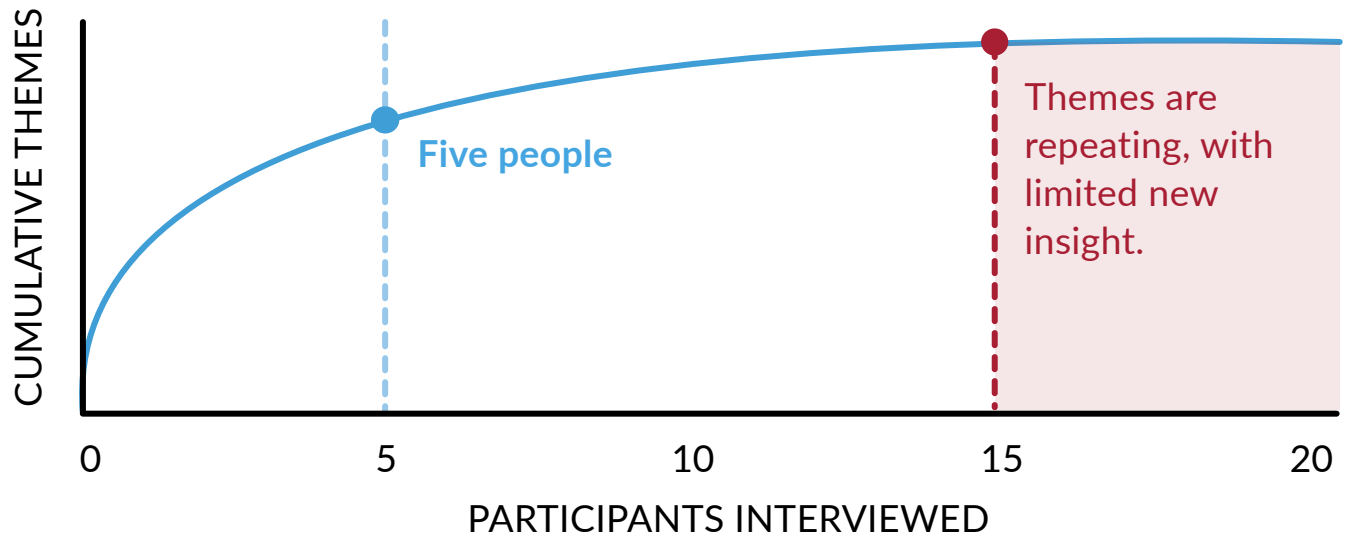

Things to keep in mind:

Meet people where they are.
Go onsite to observe clinicians, patients, or staff in the environment that they will encounter your nudge. This will allow you to understand the context you're designing for.

Focus on actions and problems.
Pay attention to what users do and what they're struggling with.

Ask open-ended questions.
Create space for people to give rich, detailed responses. Avoid any questions that can be answered by a yes or no.

Take notes by hand.
This filters out unnecessary data and helps your own synthesis. If it's hard to ask questions and take notes, ask a colleague to be your notetaker. This also prevents accidental HIPAA violations.

Avoid validation.
The goal of discovery research is not to validate your ideas. It's to gain understanding. Never ask if they like your idea.

Dig deeper.
Pay close attention to responses and listen closely for opportunities for follow-up questions. Ask "why?" to uncover root causes behind behavior.

Avoid feedback from friends or colleagues.
This includes people who report to you, like trainees and residents. They're too familiar to give honest feedback.

Get a 360 view.
Talk to patients, clinicians, staff. There might be more people than you expect who are impacted by your nudge.

Steps for Interviews and Observations

1. **Plan your conversation.** Start with questions connected to your learning goals. Then order questions from easy (their role, current state) to complex (workflow, pain points, what would make it easier to perform the behavior).

2. **Choose people to interview and shadow.** Pick people who have relevant experience to your nudge. However, avoid only speaking to people who are highly motivated to do the right behavior or very tech savvy. You want to find people who match the population of potential users, not the outliers.

3. **Ease into it.** Introduce yourself and the project. Explain your role and clarify there are no right or wrong answers. If you are recording, ask for permission to do so.

4. **Ask questions.** Begin your sessions with open-ended questions to understand their actions, habits, and pain points.

5. **Watch their workflow.** Ask if users can show you how they perform relevant tasks. Before moving into observation, signal the transition clearly — let users know you'll be watching them work and may interrupt to ask questions.

6. **(Optional) show a mock-up.** If you have a rough concept to share, like a sketch or a written description of an idea, wait until the end to share. This prevents priming them. Observe how people react to your design — don't explain or defend it.

7. **Wrap up.** Thank the participant and leave space for any final thoughts.

Synthesizing Your Data

Next, you'll want to draft a summary. Right after a live session, re-group with your colleagues and jot down notes together, in a single document. Outline any repeat patterns you observed from users.

Write your summary right after sessions — the longer you wait, the less details you will remember. Doing this as a group helps collect different information that might be easy for one person to miss.

Once you have a summary, dig deeper into your notes and focus on answers and observations related to the users' behavior. Tools like Mural or Miro can help by adding stickies with findings and grouping themes together. If you prefer working with written text, type up notes for each session and add them to a single document. Group chunks of similar feedback together to get a rough idea of themes. If you need to report findings to important decision-makers, like clinical governance groups, include direct quotes from your interviews and observations.

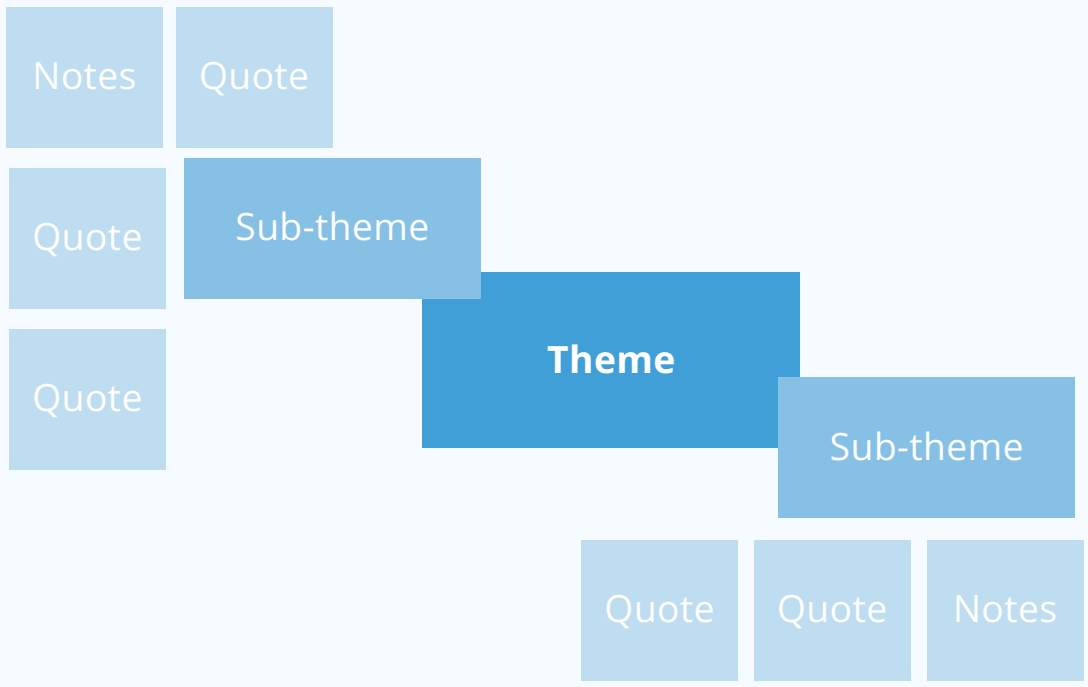

Identifying Behavioral Barriers

As you sort through your themes, focus on responses that relate to the challenges. These can surface what gets in the way of the behavior. Common signs of behavioral barriers in health care are lack of knowledge, role clarity, or insufficient time. Similarly, look closely for a user's goals and desired outcomes. These can tip you off to their motivations. Once we narrow in on the behavioral drivers, we can design a nudge that is most effective. We'll go into more detail in the next chapter.

AI in Qualitative Analysis

When faced with a large amount of data to analyze, it's tempting to turn to large language models (LLMs) for synthesis. However, the purpose of discovery research isn't just creating insights. It's to immerse yourself in the users' lived experience, deepen your understanding of their needs, and empathize with them. These are critical steps that shape how you design, pilot, and scale your nudge.

> **USE AI AS A SUPPLEMENT, NEVER AS SOLE METHOD**
>
> Literature about AI tools for qualitative synthesis shows it can give a quick summary or help speed up coding. However, current models perform poorly in identifying relevant quotes accurately, understanding cultural and emotional context, and capturing domain-specific detail. For specific tasks, AI tools can help your synthesis process but should not be the only method used during analysis.

4

Understanding and Changing Behavior

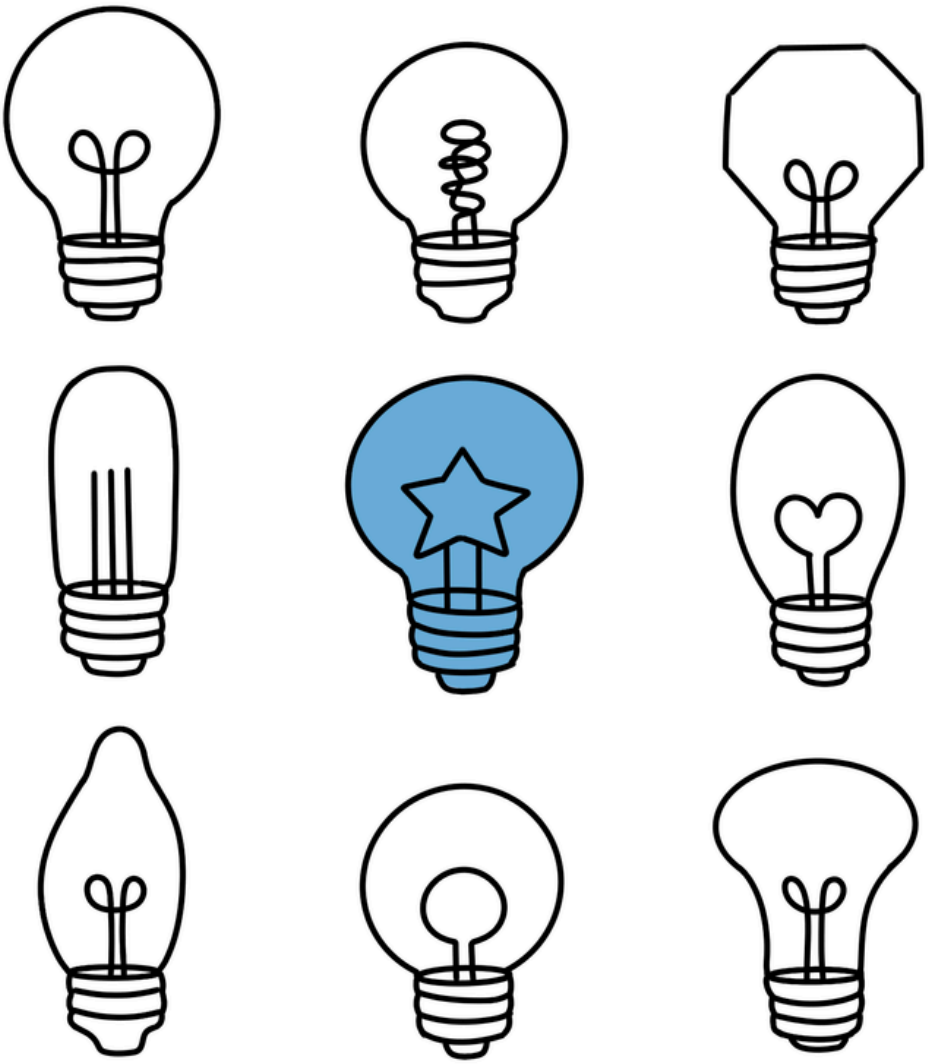

As you move from discovery to solutions, it's important to know a bit about the science of behavior change.

The COM-B model sheds light on why people aren't doing a healthy behavior—and points to what can be done to nudge them. The **B** stands for behavior. The rest of the letters stand for:

Capability
Are they physically or mentally able to do the behavior?

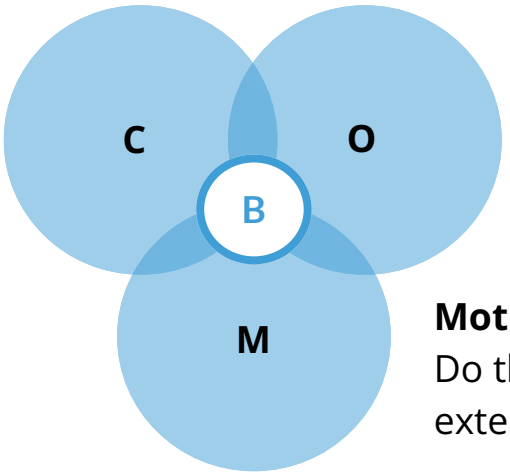

Opportunity
Do they have the time, support, or resources?

Motivation
Do they, for personal or external reasons, want to?

If a person is willing and able, and the opportunity arises, they will do the behavior. If a person isn't doing the behavior, then they're missing at least one of these ingredients.

Using COM-B

First, review the themes from your interviews and observations. What got in the way of doing the behavior? What might make it easier?

With this understanding, you can design a nudge that gives people what they need. A few tips:

If people are lacking...	Then your nudge should...
Capability	Focus less on building them up to do the behavior, and more on making the behavior itself *easier*.
Opportunity	Come at the *right time*, tied to existing workflows or habits.
Motivation	Stress *lasting* reasons to do the behavior that they care about.

Nudging for Capability

Make the behavior easier

Sure, you can train and pump people up to do a big, scary behavior. But often, it's better to focus on making the behavior itself easier:

- Give simple guides and checklists.
- Automate or remove steps.
- Set the behavior as default—if they don't opt out, it will happen. Note: Only do this for widely accepted behaviors, with little downside.

Defaulting opioid prescriptions to a lower amount

DISCHARGE PRESCRIPTION

- (•) Percocet 10 tablets
- ( ) Percocet 20 tablets

Kickstart confidence

Sometimes a person can do the behavior but lacks confidence. You can give their confidence a lift by:

- Asking them to recall when they did well at a related behavior.
- Giving examples of similar others who do the behavior.
- Encouraging and celebrating baby steps.

Texting to celebrate a small win

Today 9:00 PM

You took 5,638 steps today—way to go!

Boost their brain

Teaching and training aren't themselves nudges, but they are often needed to build capability. The problem? Knowledge and skills fade with time. Your nudge can act as a "booster" by reminding people of what they need to know:

- Give 1-3 key pointers.
- Do this before they need to act.

Naming the preferred antibiotic in the order itself

PRE-SURGERY ANTIBIOTIC

- ( ) cefazolin - preferred
- ( ) vancomycin - cefazolin allergy or SCAR to beta-lactams

Nudging for Opportunity

Give a timely prompt

Often people have time to do the behavior—just not now. You can design a nudge to prompt them at the right place, right time:

- When they have a moment
- When they are open to change
- At a key step in the workflow
- When a decision is made
- When the rate of behavior starts to drop

Prompting a naltrexone order the moment alcohol use is reported

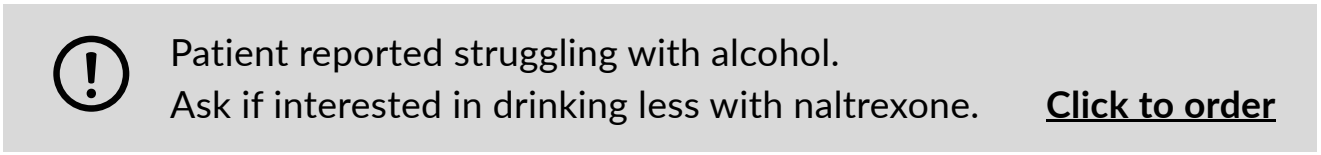

Create a supportive environment

It's much easier to swim with the current. Look for ways to change their environment so that it supports the behavior:

- Change the layout on their screen.
- Simplify the workflow or choices.
- Rearrange the physical setting.
- Show that their peers do the behavior.
- Give an authority's stamp of approval.

Showing how their prescribing compares to peers

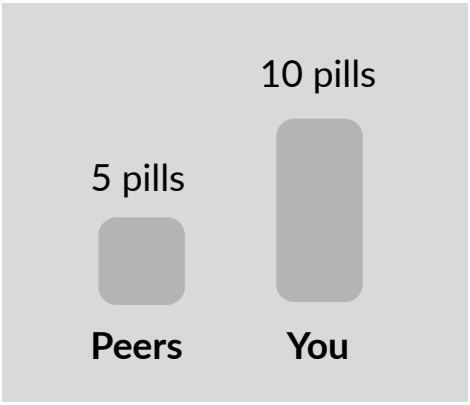

Carve out time and resources

Providing more time and resources goes beyond nudging but is sometimes needed to pave the way for the behavior. That said, there are low-touch ways your nudge can make time and resources almost at no cost:

- Block off time with calendar holds.
- Take other things off their plate.
- Offer a helpline.

Nudging for Motivation

Gently persuade

If someone is against doing the behavior, a nudge won't change their mind. Persuasion is a process. It takes listening to and answering the person's concerns. If you've done this work, your nudge can do the rest:

- Appeal to their heart and mind with a vivid story and compelling stat.
- Let them make the case: ask why they believe the behavior is important.

Frame your message

When nudging, you have a choice. You can either stress what a person stands to gain by doing the behavior or what they could lose by not doing it. Which is more motivating? Depends on the type of behavior:

- If it will help prevent illness, stress what they could gain by doing it.
- If it could detect illness, stress what they could lose by not doing it.

Gain framing

Statins improve heart health

Loss framing

If no screening, cancer may hide

Get them to commit

When a person promises to do something, they're more likely to do it. You can up the chances they'll commit by:

- Keeping the ask reasonable.
- Letting them set their own goals.

You can up the chances they'll follow through by:

- Making it a public commitment.
- Limiting future choices by agreement.

Having a clinician sign a goal their team can see

MY COMMITMENT

I will ----------

---

signed

Visible to my team

Make it a game

You can "gamify" behavior just like the most popular apps do. Because play is better than work. To gamify your nudge, recognize when people succeed at the behavior by:

- Tracking streaks.
- Awarding points, badges, statuses.
- Arranging a competition.

Tracking a daily streak of the behavior

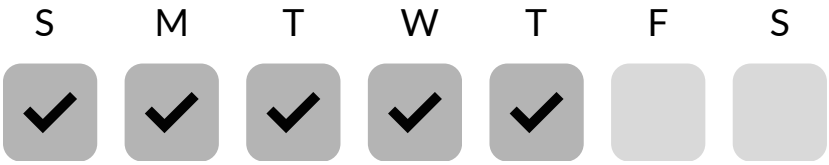

Turbocharge rewards

Offering money or threatening punishment are not nudges but can get people to do the behavior. Keep in mind: Money-based rewards don't build—and can even harm—personal motivation. And people dislike punishment, which can cause it to backfire. If you do offer rewards, get more bang for your buck by:

- Giving smaller rewards sooner, not one big reward at the end.
- Framing rewards as theirs to lose. No one likes to part with what's already theirs.

Beware of Biases

When you talk to people about why they aren't doing the behavior, you may sense there's more to the story. That's because there is! People are often steered by biases they aren't aware of.

If you can identify the biases steering people away from a behavior, you can design your nudge to steer them back to it.

Common Biases and What to Do About Them

Bias	Definition	Example	How to Fight
Who's responsible?	When everyone can take charge, nobody does	Staff, nurses, and doctors can order an interpreter, but none do	Make roles clear; assign responsibility
Status quo	Liking what's familiar, fearing the downsides of change	Pushback on changing the criteria for a drug test	Highlight benefits of change, costs of status quo; prompt new vs. old choice; make new behavior recommended default option
Present focus	Caring most about immediate rewards and consequences	Giving less effective antibiotic for fear of allergic reaction to most effective	Reward now what's best for future
Stereotyping	A person's race, sex, etc. colors how we see them	White nurses order drug test more often for Black patients	Standardize how decisions are made; remove opinion
Ostrich	Avoiding important news that could be bad	Patients with symptoms put off getting screened	Remind person about what won't change, and their options if news is bad
Top of mind	Recent, vivid, or emotional info wins out	Reading horror stories stops folks from getting vaccine	Make info supporting best practice recent, vivid, or emotional

5
Designing for Scale

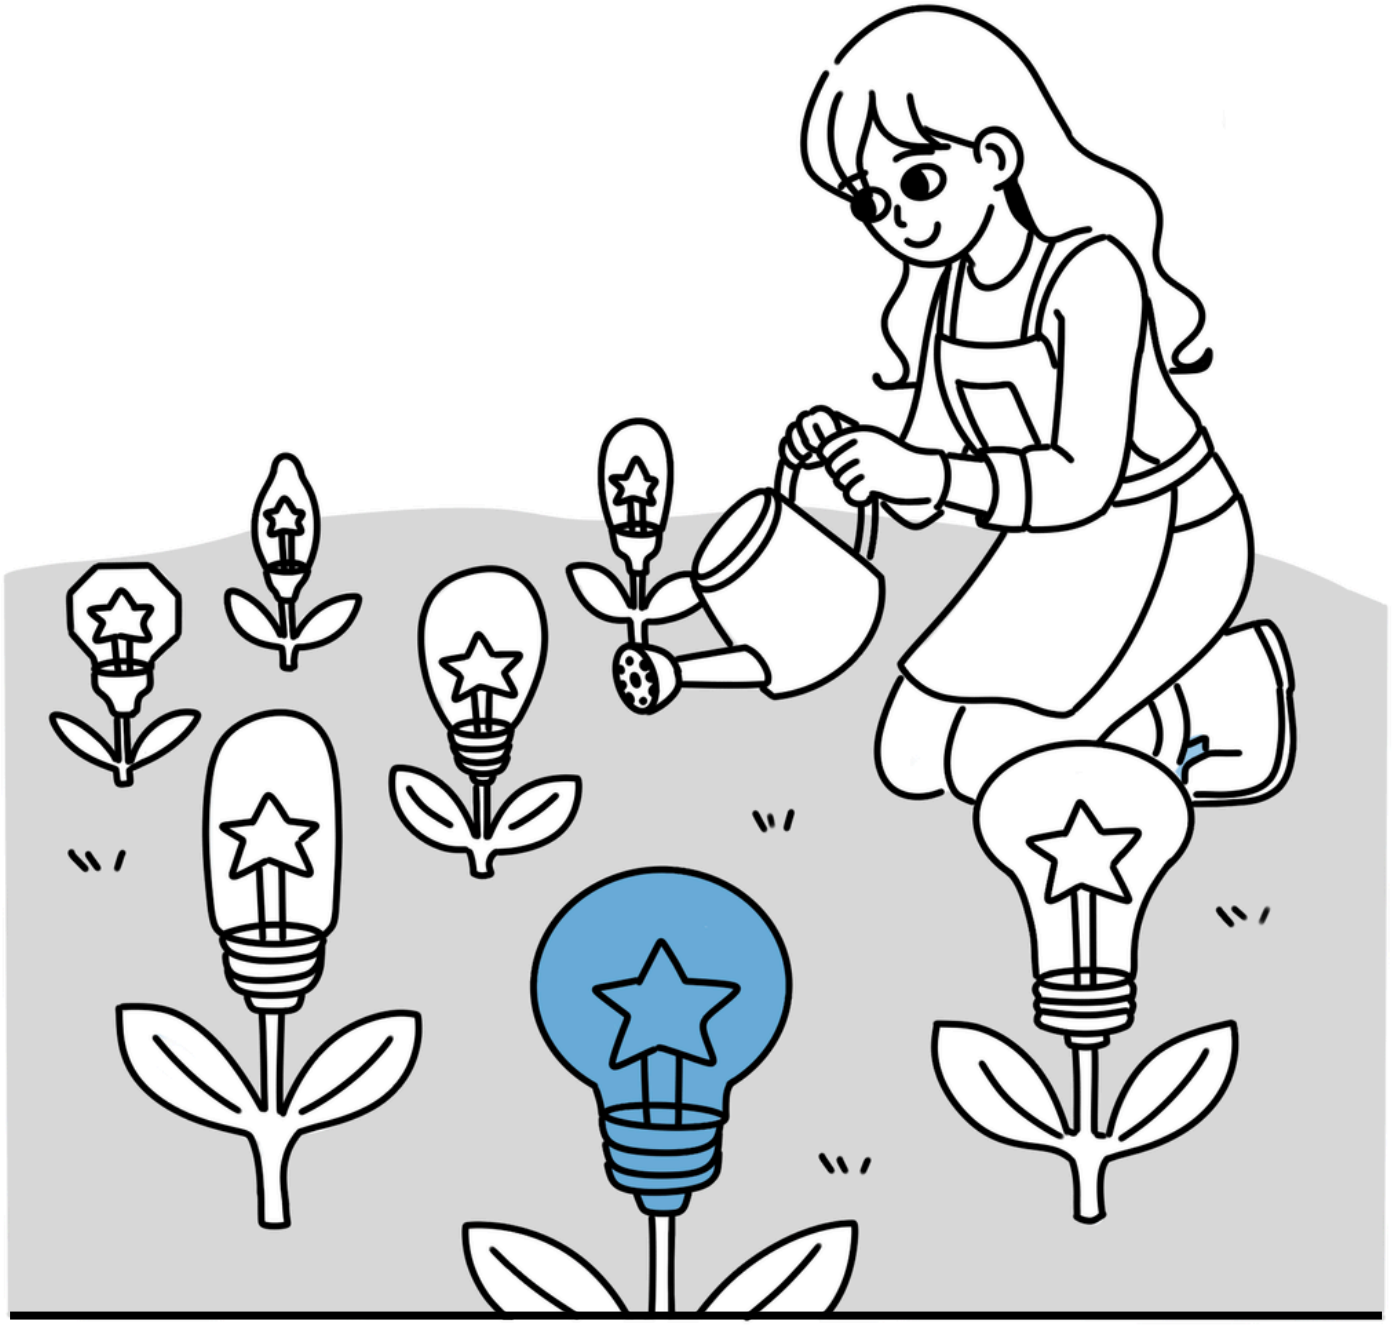

Now that you've identified the problem and are exploring potential solutions, consider how to make your solution more likely to scale.

Spread vs. Scale

Spreading means taking a proven nudge from one setting to another and using it in the same way. It assumes the nudge will work just as well in the new setting without changes. That may be true when the new setting is very similar, like another care team in the same practice or department. But spreading often does not work as well when the nudge moves to a different type of care setting, or department.

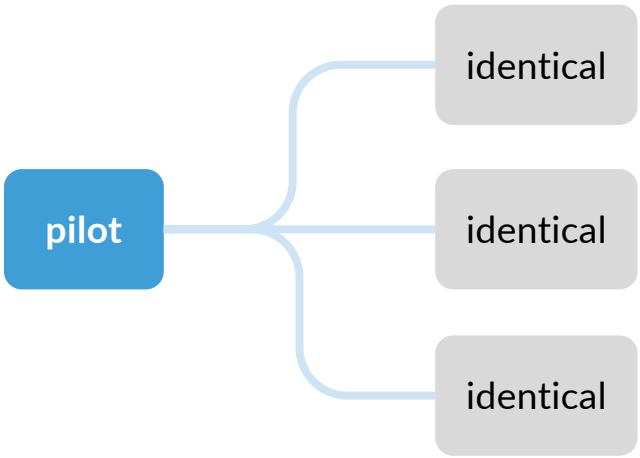

Assumes the new setting is the same

Scaling means expanding a proven nudge to a new care setting and making changes, so it fits the new setting. This includes considering staffing, workflows, technology, leadership priorities, team norms, and patient needs. The goal is to ensure the nudge will be as effective in the new setting. Scaling should be planned carefully so the nudge is accepted, set up for success, and designed to last.

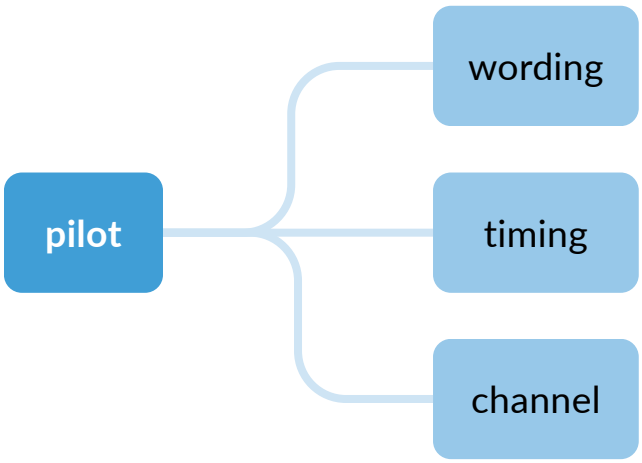

Core mechanism is same, other components customized

Importance

Scaling helps a successful pilot reach more people and have a bigger impact. Planning for scale early helps make sure a nudge won't only work in one place but can work well in many different care delivery settings. Pilots often work because they get extra attention, and have strong champions and short timeframes. Designing for scale helps you:

Identify system needs early. Figure out what the nudge depends on—EHR (electronic health record) constraints, workflows, staffing, leadership support, or incentives before these become problems.

Make better design choices. When planning for scale, use simple defaults, clear decision points, and nudges that fit naturally into current workflows.

Avoid expensive redesigns. If scale is not considered during the pilot, the nudge may need major changes later—or may not work when expanded.

Increase your potential reach. A pilot shows that a nudge can work in one place. Scaling tests whether it can keep working across different settings over time.

Plan for long-term success. Clear ownership, regular monitoring, and ongoing support make it more likely the nudge will keep working after the pilot ends.

A pilot should not only ask, "Does this nudge work here?" It should also ask, "What would need to be true for this nudge to work in other places?"

Design for Scale Checklist

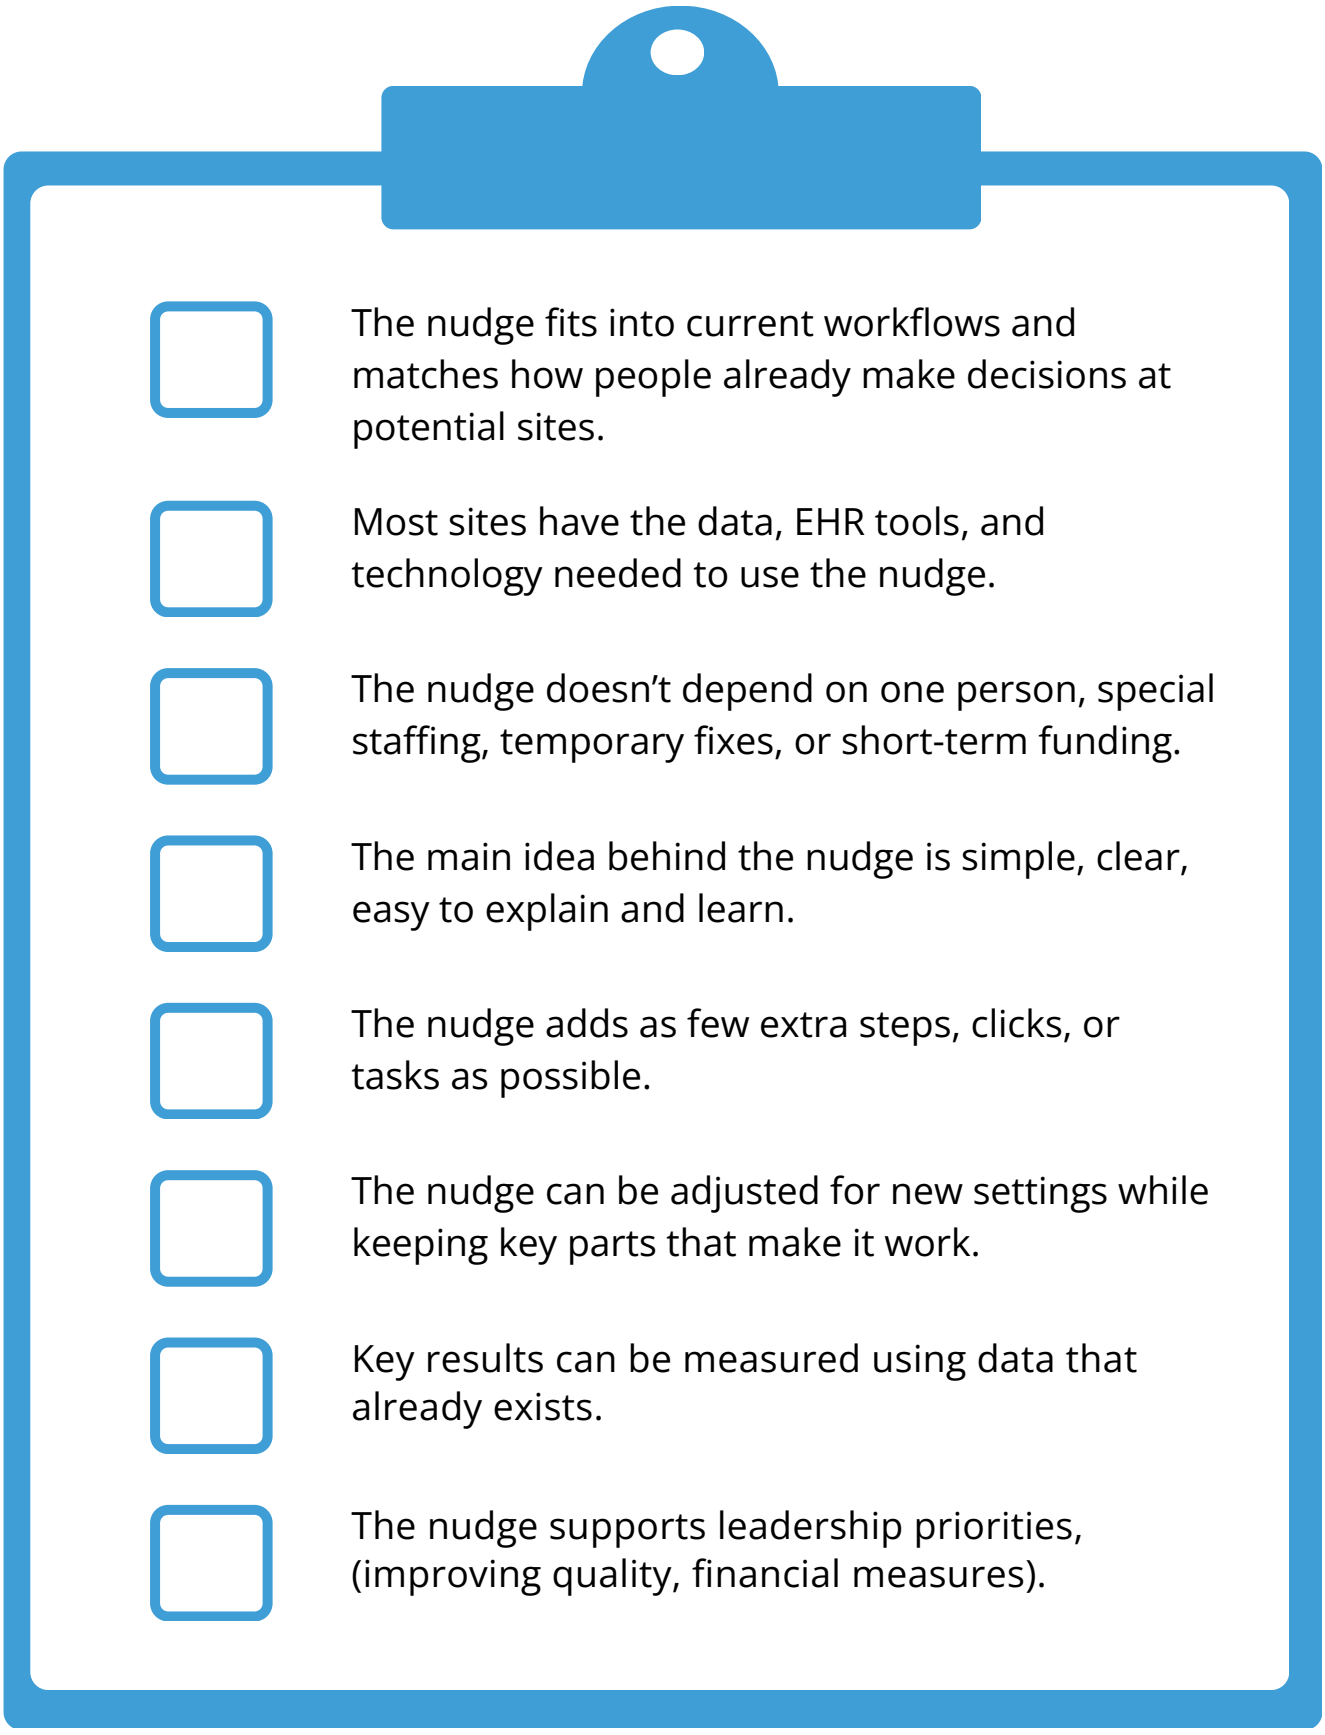

- [ ] The nudge fits into current workflows and matches how people already make decisions at potential sites.
- [ ] Most sites have the data, EHR tools, and technology needed to use the nudge.
- [ ] The nudge doesn't depend on one person, special staffing, temporary fixes, or short-term funding.
- [ ] The main idea behind the nudge is simple, clear, easy to explain and learn.
- [ ] The nudge adds as few extra steps, clicks, or tasks as possible.
- [ ] The nudge can be adjusted for new settings while keeping key parts that make it work.
- [ ] Key results can be measured using data that already exists.
- [ ] The nudge supports leadership priorities, (improving quality, financial measures).

6
Prototyping

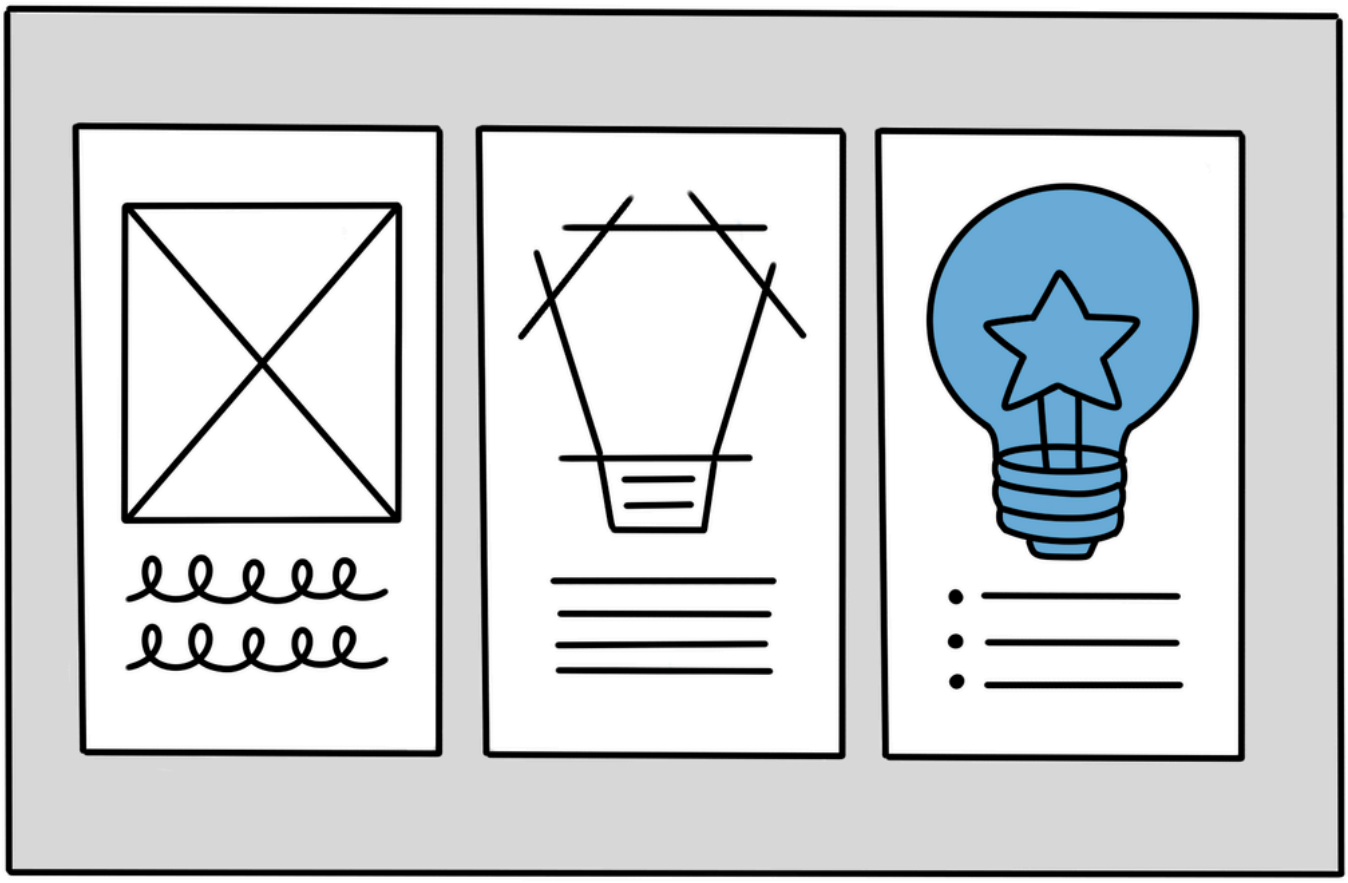

Now that you have narrowed your focus to potential scalable solutions, it's time to make them feel real. Prototyping means building a simple version of your idea to see how people react to it. It's a way to learn quickly.

There are many ways of prototyping. People can create realistic, interactive digital prototypes to test with users. For most people, lo-fi prototyping will be the most accessible way to create mock-ups of your nudge.

Benefits of Lo-Fi Prototypes

Help everyone understand.
A real, visual, prototype communicates an idea more clearly than words alone, reducing misunderstanding.

Explore multiple concepts.
Expand your thinking by creating several versions of a solution. This helps you catch problems early, before they become bigger issues.

Reduce risk by testing early.
Test out your ideas without investing too much time or resources.

Gather user feedback.
Put concepts in front of users early so you can iterate and improve it before building your nudge.

Works for any type of solution.
Whether physical or digital, the prototyping process is the same.

Paper Prototypes

Paper prototyping is when teams create a paper model of an idea to visualize a concept. They are most useful when developing early concepts to quickly test ideas, layout, and workflows at a low cost. Paper prototypes are useful for trying out new ideas. They help teams work together and test quickly, before investing in anything digital. To make a paper prototype of a digital nudge, you can print pictures of the EHR screens or draw basic shapes to get your idea across with users.

When designing solutions at The Nudge Unit, we've used paper prototypes extensively. For a project redesigning a labor floor urine drug screen order, we created a paper mock-up of an order requiring clinicians to select which guideline was met to order the test. You don't need design training to do this: all we used was a screenshot and PowerPoint. For another project, we created a paper prototype of a EHR banner encouraging clinicians to use certified interpreters for patients who do not speak English. We cut the paper and taped the banner mockup to their computer screens, observed how they interacted with it, and asked for feedback. In both cases, developing the paper prototype helped align the project team and validate an early concept, before we invested more time and resources in the concepts.

Example of EHR-based banner.

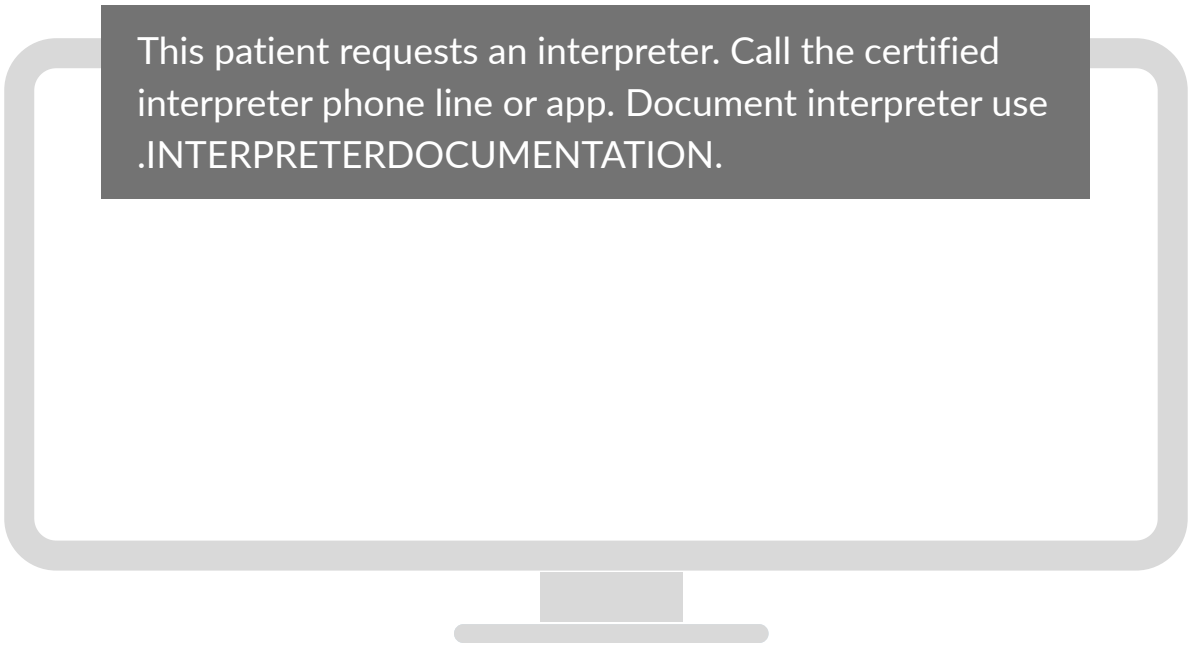

Once you've created a prototype, it's time to get it in front of potential users. Recruit five potential users to test your design. If it's complex, aim for slightly more users, six to eight participants.

NEW THEMES DISCOVERED PER ADDITIONAL PARTICIPANT

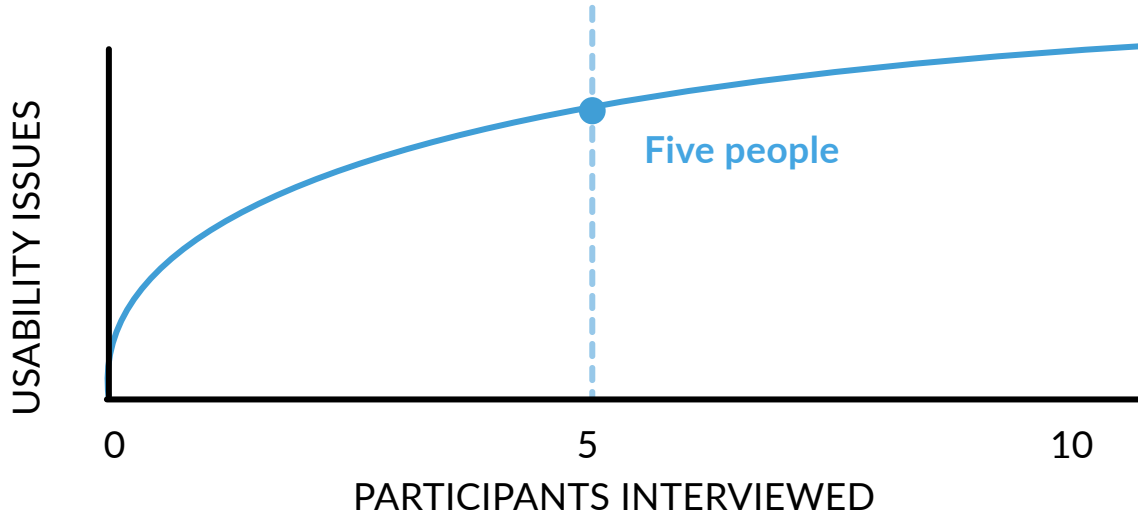

Instructions for user testing:

Ask open-ended questions

This encourages users to give rich, detailed responses and avoid yes/no responses. Ex. "What are your thoughts on the new order?"

Use the "Think Aloud" method

Encourage people to talk about their thoughts as they interact with the prototype. Avoid giving direct answers or hints. Ex. "Can you think out loud as you go explore the prototype?"

Step back

Your role during user testing is to observe. Avoid giving participants direct answers or hints. If they ask you questions, ask the question back to them: "What do you think this feature is supposed to do?"

Allow struggles

If a person gets stuck, don't rush to help them. Give them time to figure it out, as this will help identify issues. If they seem completely stuck, use open-ended prompts like: "What are you thinking at this moment?"

7
Evaluation Plan

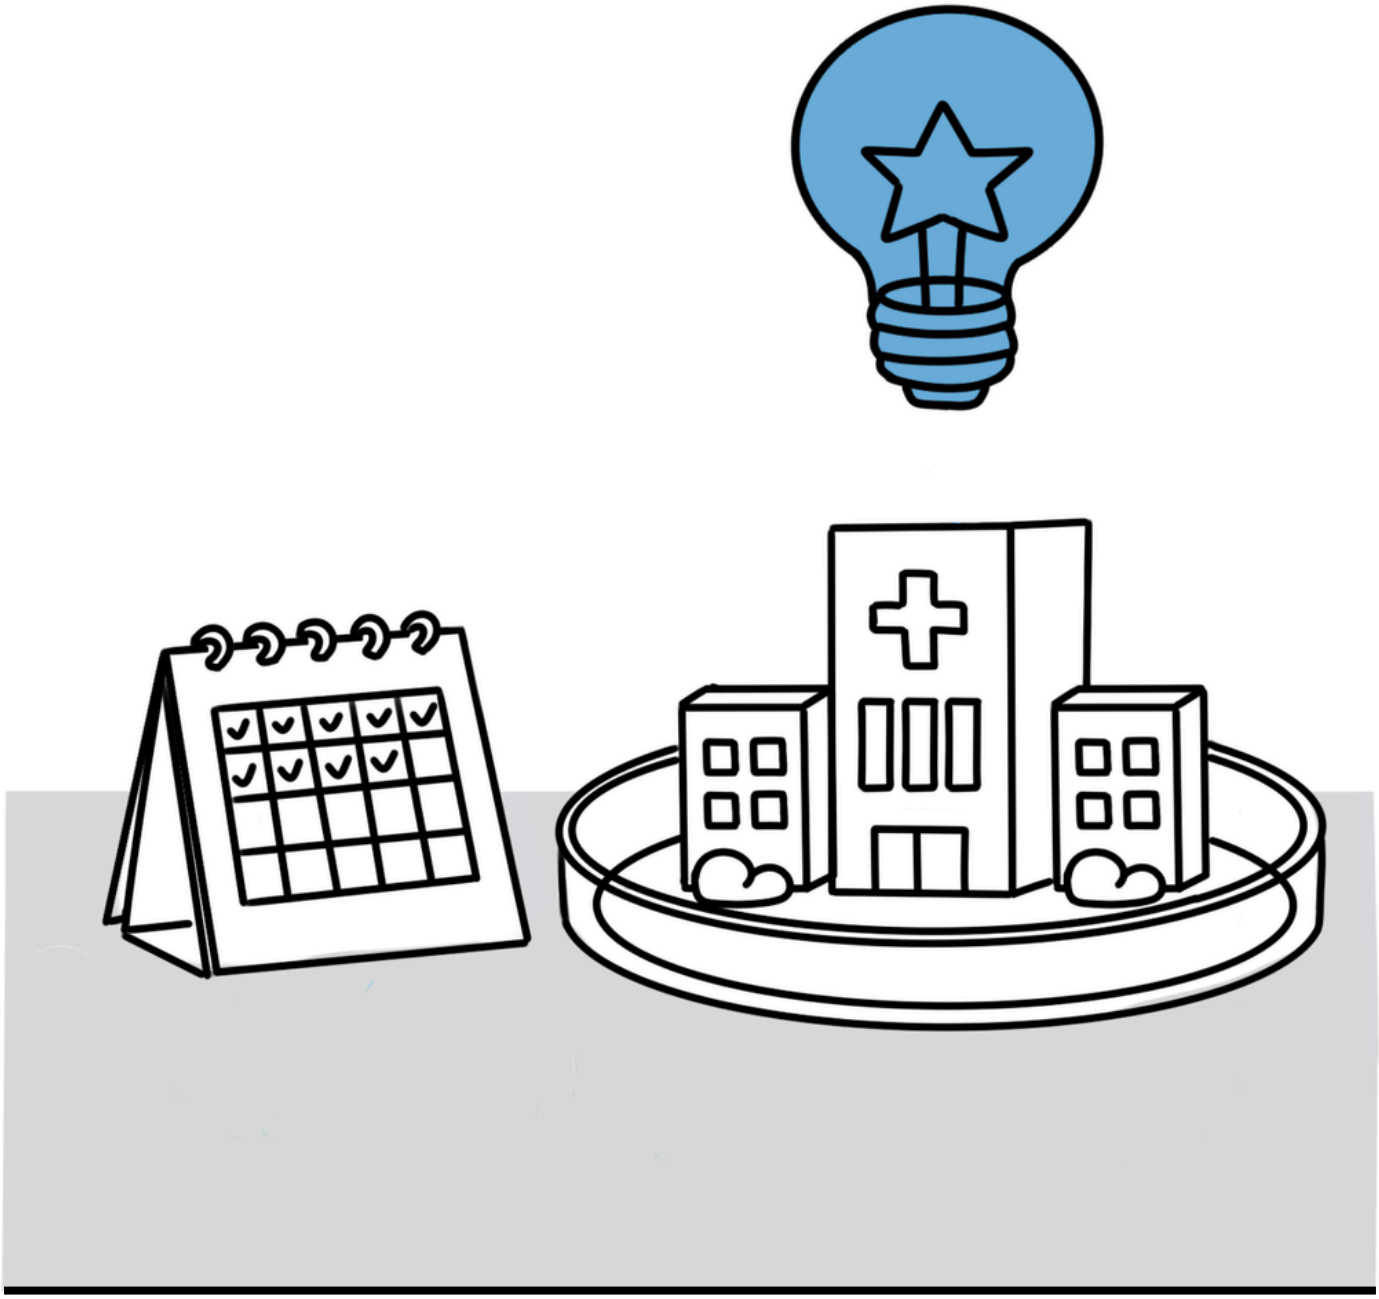

Before launching your nudge, you need a plan to evaluate its success. This means choosing your outcomes and how you'll analyze them.

Choose Your Outcomes

Imagine a nudge that sets the default statin prescription to be a 90-day supply instead of 30-day. What should you measure to understand its success?

North Star

Begin by naming your North Star. This is the long-term goal of your nudge, why the work matters. For the statin nudge, your North Star is fewer heart attacks and deaths. You can certainly measure your North Star outcome. But often it will be too rare or too far in the future to gauge if your nudge is working. Instead, you need a good primary outcome.

Primary Outcome

This is the near-term target of your nudge, what you hope to see a change in. A good primary outcome is:

- Believed to affect your North Star
- Easy to measure
- Likely to be moved by the nudge
- Meaningful to patients and/or staff

Common primary outcomes for health care nudges include number of medication orders, percent of patients screened, and length of stay. For the statin nudge, your primary outcome is percent of statin prescriptions that are 90-day.

Secondary Outcomes

Your nudge will likely impact other things worth measuring. Secondary outcomes include:

- Other steps in the causal chain from nudge to North Star (% of patients taking statin at 90 days, "bad" cholesterol levels)
- Patient reports of how they're doing (symptoms, treatment satisfaction, reasons for stopping statin)
- Unintended consequences (pharmacy supply issues, lower prescription fill rate among low-income patients)

Implementation Outcomes

Say your nudge doesn't move the needle on your primary outcome. Is this because the nudge doesn't work, or because it wasn't carried out right?

To figure out what's happening and what you can do to improve your nudge, it's important to measure:

- What percent of sites or staff adopted the nudge
- What percent of targeted people were reached by the nudge
- How consistently the nudge happened as planned
- Whether the nudge was maintained over time

NORTH STAR *Why the work matters*	*fewer heart attacks and deaths*
PRIMARY OUTCOME *What you hope to move*	*% of statin scripts that are 90-day*
SECONDARY OUTCOMES *Everything else worth watching*	*adherence at 90 days, unintended effects*
IMPLEMENTATION OUTCOMES *Was the nudge delivered right*	*adoption, reach, fidelity*

Choose Your Analysis

Your analysis will reveal if your nudge moved the primary outcome. Your options for analysis are tied to your study design. We'll look at three options, from least to most rigorous.

Pre/Post

Measure the outcome before the nudge. Then, turn on the nudge and measure it again.

- **Pros:** Fast, needs no comparison group, and uses data you likely already collect.
- **Cons:** Can't know for sure if your nudge affected the outcome. Maybe something else changed around the same time, or maybe the outcome was already trending up or down.

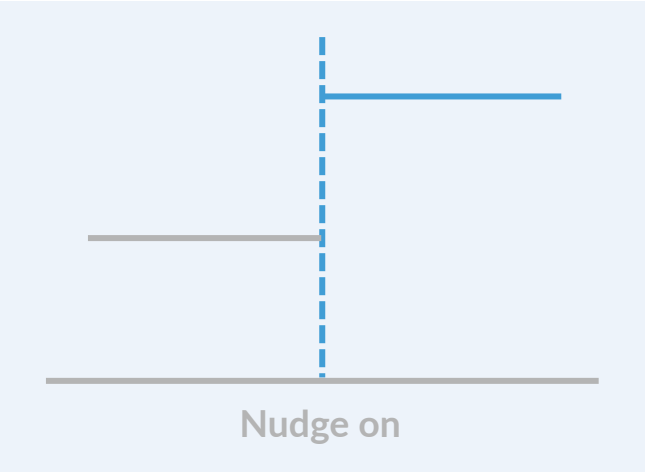

- **Ways to strengthen:**
  - Use many time points so you can see if the nudge led to a sudden jump or a change in slope, and what happened when the novelty wore off.
  - Before you launch, decide the amount of time and the number you need to hit for success and put it down in writing.
  - Track a "control outcome" that the nudge should not affect. If it moves too, something other than your nudge may be driving the change.
  - Keep a dated log of every other change to the workflow during the study period.

Difference-in-Differences

Compare the change over time in a group that got the nudge to the change over time in a similar group that did not.

- **Pros:** Useful when you can't randomize but have a similar comparison group (another clinic, division, unit, or hospital) and data from before and after the nudge went live. Controls for seasonality, new policies, staffing changes, EHR upgrades, etc.
- **Cons:** Doesn't control for every difference between the two groups.

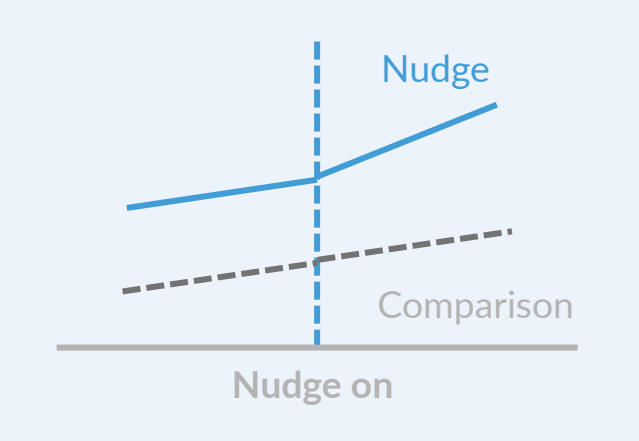

- **Ways to strengthen:**
  - Ensure that the trendlines for the two groups were parallel before the nudge.
  - Choose a comparison group exposed to the same system-wide pressures, but not the nudge itself.
  - Analyze at least 6-12 months of data before and after the nudge.

Randomized Controlled Trial

Randomly assign the nudge, then compare the outcomes of those who did and didn't get it.

- **Pros:** Considered the "gold standard"; lets you confidently say that the nudge caused any changes in outcomes.
- **Cons:** Hard to pull off when nudging in a health care setting.

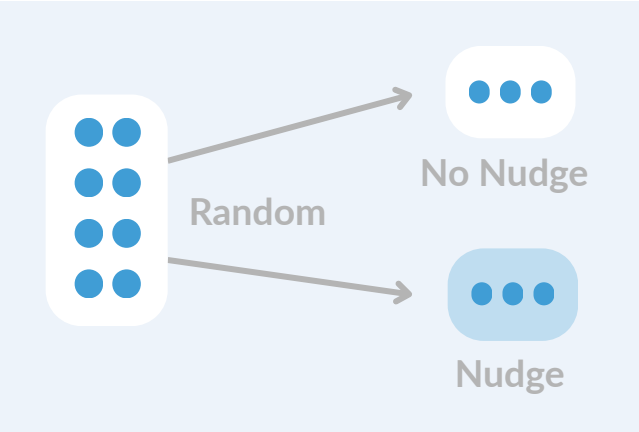

- **Ways to strengthen:**
  - Randomization can occur at many levels: patient, clinician, clinic, department, or hospital. Choose a higher level (e.g., clinic) if you're worried about individuals randomly assigned to a nudge (e.g., clinicians) comparing notes with those in the control group.
  - Decide in advance if there are any characteristics (e.g., patient sex, clinician tenure) that could affect your outcome. Get these data and ask your analyst to control for them in the analysis.
  - Decide in advance how much time should pass or how many randomly assigned cases there should be to evaluate the primary outcome.

Choosing among the three: There is a clear tradeoff between rigor and ease. The right choice is the most rigorous design your setting will support.

Checking for Equity

Certain groups experience worse health outcomes. A nudge can improve outcomes overall while leaving disparities untouched. When evaluating a nudge, it is important to check if it has narrowed or—more concerningly—widened existing gaps.

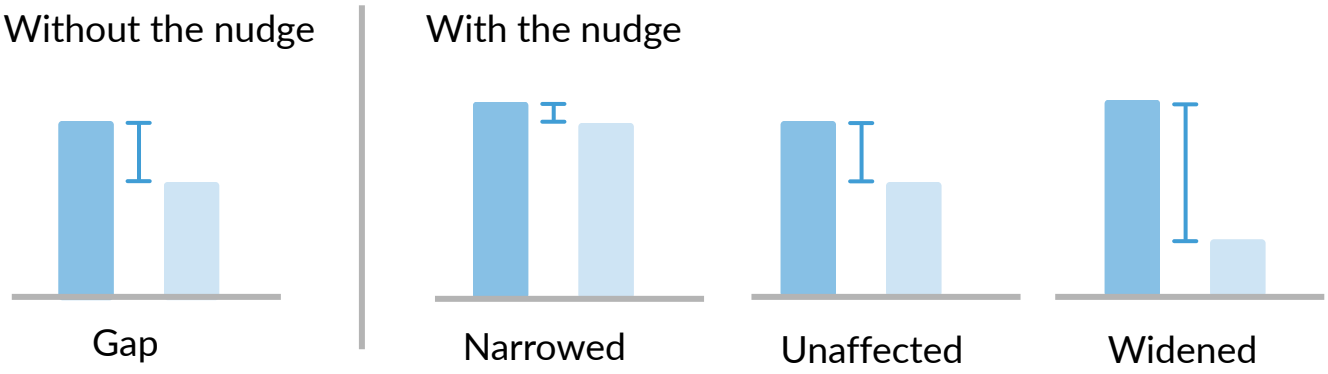

Let's say that for your primary outcome, there is reason to suspect a difference between those who do and do not have private health insurance. You'll want to compare the size of this gap without the nudge to the size of the gap with the nudge.

How Much Data Do You Need?

You may be limited to what data you have access to or other external constraints. However, a power calculation is the best way to know you have a large enough sample size to prove your nudge's effectiveness. A power calculation answers one question: if this nudge works, is our study big enough (e.g. enough power) to notice? Put another way, if you have a large sample, you will have enough variation in the group that we can say with confidence the thing causing the change is the nudge (and not unrelated factors). Power is the probability of detecting a real effect. Best practice is to aim for 80% power, or an 80% chance of finding the effect.

Four important numbers:

1. The **baseline** of where the outcome sits today (e.g., 40% of patients get the test).
2. The **smallest meaningful change** worth detecting, not the change you hope for, but the smallest improvement that would actually change a decision.
3. The **power percentage** (which is typically 80%).
4. The **significance level** (which is typically 0.05).

If the calculation says you are underpowered, the options are to run it for longer, add sites, raise the effect size you are willing to call meaningful, or reframe it as a "learning pilot" and report it descriptively. The learning pilot can help you understand and assess feasibility.

Rules of thumb worth knowing:

- Rare outcomes need far more patients than common ones. If your outcome occurs in 2% of patients, expect you will need a large sample size.
- Continuous outcomes (length of stay in days, minutes to first order) are more statistically efficient than yes/no outcomes.
- Build in attrition and eligibility loss. Exclusions and patients who never trigger the nudge will shrink your usable sample, often more than you expect.

> **WHEN TO BRING IN A STATISTICIAN**
>
> Free online calculators can handle standard cases. Involve a statistician or analyst when you are randomizing groups rather than individuals or running a more complex study design.

Timing of Evaluation

Decide these checkpoints, and what you will do at each one, before the nudge goes live.

Baseline
Establish the benchmark, confirm the data pull works end-to-end, and look at the outcome by patient population to identify equity gaps the nudge should narrow.

1

2

Pilot (first days to weeks)
Test on a small sample. The only question here is whether the nudge is technically delivering and working or firing for the right patients, at the right moment.

Early post-launch (2–4 weeks)
Watch for a novelty spike and any unintended consequences (ex. alert volume, override rates, clinician complaints).

3

4

Midpoint
Check that recruitment is on track against the power calculation, and that data collection is still complete.

End of study (pre-specified)
Interpret the primary outcome, secondary outcomes, and any unintended effects, then decide whether to scale, adapt, or stop.

5

6

Sustainment (6–12 months after)
Re-check the outcome. Nudge effects can regress as the prompt becomes background noise.

Source of Data

Confirm you can get every data element you need and that you have access to it in real-time. Discovering a gap after going live usually means the baseline period is unrecoverable.

EHR structured data: e.g. orders, diagnoses, medications, labs, vitals, encounters, and scheduling data. The workhorse for most outcomes. Know the current state of research data access at your organization and what approvals it requires.

Nudge exposure and interaction logs: whether the alert fired, who saw it, and what they did (accepted, dismissed, overrode, and the reason).

Data collected outside the EHR: surveys, brief interviews, or focus groups with clinicians and patients can be very helpful. These qualitative insights can provide the "why" the nudge was noticed, understood, trusted, and workable.

Operational data: census, staffing ratios, visit volume. Not usually the outcome, but the explanation when results move for reasons unrelated to the nudge.

8
Build and Scale

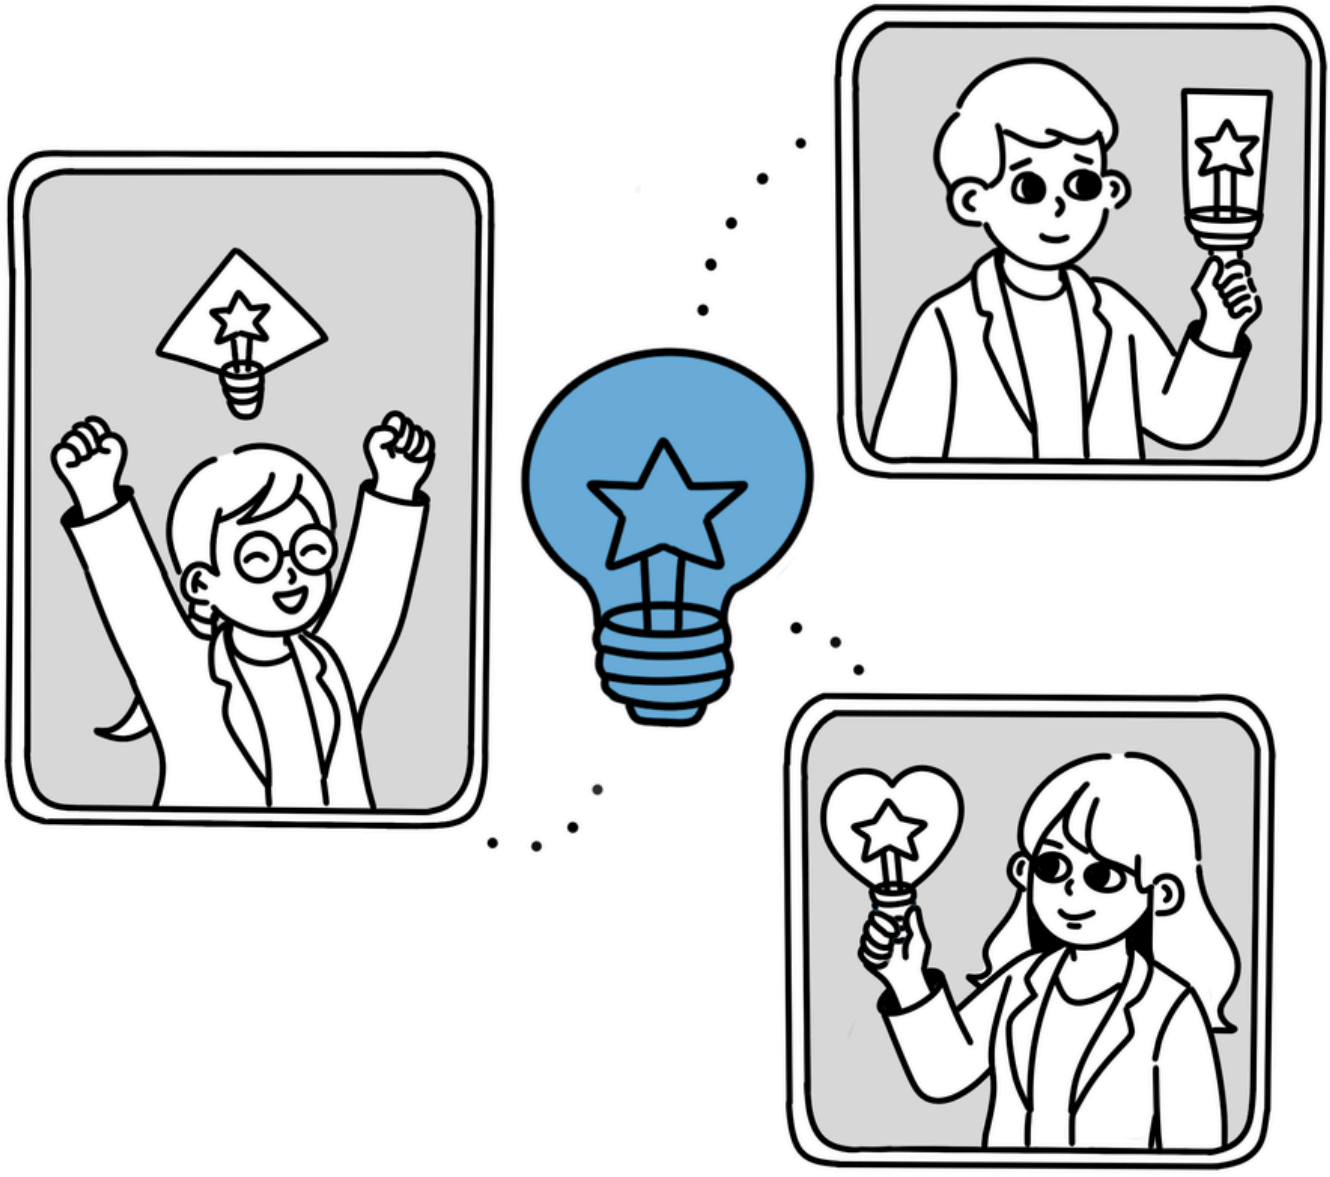

Getting Approval

One of the most crucial steps to getting your nudge into the world is getting approval from relevant clinical governance groups. Not only does this give you permission to implement your idea, but having a stamp of approval from authority will make it easier to get other clinicians on board. Whether you're testing your nudge in a pilot or making a change for quality improvement, make sure you talk to whoever leads the department at the very least. If your nudge changes something in the EHR, you may also need to talk to an informatics governance group.

Making It Real

The last thing getting approval does is tell you what steps you'll need to take to get your nudge into the workflow. For digital EHR nudges, this will likely be a ticket to your IS (Information Systems) team, or whoever manages your EHR. For physical nudges (like papers taped to a door, bracelets, or magnets) you may be able to create them yourself or speak to your procurement team.

> **BRING THE EVIDENCE WITH YOU**
>
> Tip: To show how rigorously you tested your idea before seeking approval, include quotes from your user interviews and pictures of your prototype. You may also include your business case, which we'll talk about in the next section.

Scaling Your Nudge

If your nudge pilot test was successful, congratulations! The next step is scaling. **Scaling** means deliberately expanding and adapting your nudge to new departments, sites, care delivery settings, or patient populations.

To scale successfully, you need to understand what makes your nudge work, what can be changed without breaking it or making it ineffective, and what must be in place before you expand.

Assess Your Scale Readiness

A nudge that succeeded in a pilot test or at one location may fail at scale if it relies on a single champion, a workaround, a dataset, or technology available only at your site. Hopefully, you designed your nudge for scale, so you already avoided some of these obstacles. However, there are other important readiness questions that you should be able to answer "yes" to before your nudge is scale ready.

Can you answer Yes?	If No: what to do
Does your nudge avoid "unscalable ingredients" — a single champion, a custom technology, an unusual staffing situation, or time-limited funding?	*Redesign the nudge to remove those dependencies before scaling.*
Are leaders at the scale site committed to putting the nudge into practice? (includes governance approvals, identifying a maintenance and cost plan)	*Find a leader who can sponsor putting the nudge in place at the scale site and who can help develop a plan to maintain the nudge over time. Nudges without clear owners or operational funding drift or disappear.*
Do you have the data and tools needed to launch the nudge and track results at the new site?	*Identify data gaps early. Meet with your data and technology partners before committing to expanding.*
Was the nudge tested on a similar population of the sites you plan to scale to?	*Consider a small study to test how your nudge will be received by a new population without repeating a full pilot. Also, roll out to sites that are most like your pilot population first.*
Have you considered equity? Will scaling reach underserved populations, or risk widening existing care gaps?	*Evaluate your pilot data by race, insurance status, language, and site to see who the nudge is/is not reaching. Build equity goals into your scale plan.*
Have you defined the core components of the nudge so your scale site knows what can and cannot be changed?	*See the table below.*

Define What Must Stay the Same — and What Can Change

Scaling requires adaptation, but not all parts of your nudge can be changed. Think of your nudge like a recipe: the core ingredients (the behavioral mechanism) must stay constant, but the choice of seasoning or the type of pan used to make the recipe may be able to be adjusted to what is available or preferred locally.

Core components are the parts that, if removed or changed, would affect how the nudge works. These must stay the same. Examples include the default option in an opt-out design; loss-framing in a social norm message; the decision point in the clinical workflow where the nudge occurs.

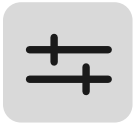

Adaptable parts of the nudge are elements that can be changed to match a local setting without causing the nudge to lose its effectiveness. They are typically related to how the nudge is presented rather than how it functions. Examples include adjusting the reading level of a nudge or translating it into another language. It could also include things like shifting the nudge from morning to afternoon to match hospital rounding patterns or using another type of text messaging software.

Significant Changes to the nudge should be considered with caution. These include changes that are more closely connected to how the nudge functions. Examples include which care team member delivers or receives the nudge, how choices are presented or triggered, the delivery method of the nudge (text vs. email message) or how often the nudge occurs. These kinds of changes should be tested and monitored to see if they are truly adaptable without making the nudge less effective.

Use the table below to evaluate any change to the nudge before making it.

Stoplight Guidance for Adaptations

Proceed	Check & Monitor	Stop & Reconsider
Language & Framing		
Adjust the language for a new audience e.g., Change wording to better fit people with lower health literacy or who speak another language.	Change who delivers the message e.g., Have a nurse deliver the message instead of a physician. Ensure message still comes from someone patients and staff trust.	Reframe a loss-framed message as gain-framed, or vice versa e.g., Changing a message that highlights a missed opportunity into one that provides information about benefits may affect how the nudge works.
Defaults & Choice Architecture		
Add defaults that support higher-risk groups e.g., Automatically schedule more frequent follow-up reminders for patients who face social or economic obstacles to care.	Switch from a default to an active choice e.g., If local policies require users to make an active choice, track whether completion rates drop	Remove the default completely e.g., Eliminating the default eliminates what made the nudge work
Timing & Channel		
Adjust when the nudge appears. e.g., Deliver the nudge at the morning huddle vs. afternoon huddle, if that is when the care team meets — if the decision point is the same.	Move to a different delivery method e.g., Use a pop-up alert versus a static alert to prompt a behavior.	Deliver the nudge after the opportunity to act is passed e.g., Send an order-at-discharge nudge after discharge orders are placed, meaning the action window has closed.
Reach & Eligibility		
Include additional groups that could benefit. e.g., Add uninsured patients who were excluded from the pilot.	Reduce how often the nudge is delivered to high-volume sites e.g., Move from daily to weekly reminders. Check that effectiveness does not decrease.	Exclude high-risk or underserved groups e.g., Limiting the nudge to commercially insured patients for convenience widens existing care gaps.

Build and Communicate the Business Case

Showing that a nudge improves clinical outcomes is important, but it's not enough to secure long-term support. Leaders and operational sponsors need to know what it will cost to expand the nudge, what benefits it will deliver, and how those costs and benefits may change as it reaches more people and care settings. A strong business case helps build leadership support and turn a successful pilot into a nudge that can be used more widely.

Health care nudges are typically low-cost relative to their impact, but it's important to make the case thoughtfully and clearly. Use the table below as a starting point for your business case.

Cost dimension	Questions to answer before you scale
Implementation costs	What will it cost to set up the nudge at each new site? Account for EHR build time, staff training, project management, and any materials or communications.
Ongoing operational costs	What does it cost to run and maintain the nudge after launch? Include data monitoring, IT support, periodic refresh, and clinical champion time.
Value generated	What is the measurable return? Quantify in clinical terms (e.g., preventable admissions, care gap closures) and, where possible, financial terms (e.g., revenue from quality bonuses, avoided costs).
Economies of scale	Does the per-site cost fall as you add sites? Identify shared resources, bulk contracting opportunities, or staffing models that reduce marginal cost.
Funding source	Who is paying and for how long? Consider start-up funding (grants, innovation budgets) and steady-state funding (operational budget lines, value-based contract dollars).

Lead with what leadership measures. Pair your clinical outcomes with the metrics your organization's leadership cares about — readmissions, length of stay, patient experience scores, or value-based contract performance. A nudge that prevents 10 unnecessary admissions tells a stronger story when you also estimate the avoided cost per case.

Distinguish launch costs from steady-state costs. Many nudges are cost-intensive up front (EHR build, stakeholder alignment, training) but cheap to run once live. Show this curve explicitly: a nudge that costs $40,000 to build but $4,000 per year to maintain looks very different from one that requires ongoing per-site investment.

Identify the path to a budget line. Grant and innovation funds get programs started but rarely sustain them. The goal is to move from a project budget to a planned operational budget line for a department or system. This requires identifying which leader has both the authority and the incentive to own the cost.

Plan for Governance, Ownership, and Sustainability

Nudges that survive at scale are almost never those with the strongest evidence — they are those with the clearest ownership and the best institutional fit. Before your go-live date at any new site, the following people should be identified by name or team.

Role	Responsible for
Clinical champion	Encouraging adoption and use of the nudge within the clinical team; raising concerns when needed; reviewing results
Data / analytics owner	Keeping data flowing correctly; running reports; identifying changes in results
EHR / IT owner	Managing the technical set up; making updates and handling version changes
Operational sponsor	Maintaining funding; approving expansion plans; resolving conflicts

Monitor, Adapt, and Evaluate

Scaling doesn't mean the work is done. After a nudge is rolled out broadly, you need to make sure it continues to work as intended. Plan how you will monitor it from the beginning.

Adoption: Are people using the nudge? Measure clicks or order completion rates.

Workflow impact: Is it making work easier or creating burden? Track time spent, clicks, and feedback from clinicians.

Results: Track at least one outcome, even if it is not perfect. Note any limitations.

Staying true to the nudge: Core components that make the nudge effective are preserved.

Adaptations: All changes to fit the local workflows, staffing, etc., are documented.

Equity: Check whether results differ across patient groups or locations. Investigate any growing gaps before expanding.

Set a re-evaluation cadence. Nudges are not “set and forget.” EHR changes, clinical guideline updates, staffing turnover, and patient population shifts can all lower effectiveness over time. We recommend a formal review at 3 and 12 months post-launch at each new site.

Conclusion

Congratulations! You've now gone through the process of creating a nudge from the original idea all the way to scaled at multiple sites. Creating impactful, evidence-based nudges can be a long process, but taking time at every step will save you wasted effort and time in the long run.

> **Improving the health care system may seem like an uphill battle, but sometimes the best place to start is with just one small nudge.**

Credits

The Penn Medicine Nudges in Health Care Toolkit was written and produced by the Penn Medicine Nudge Unit, with support from the Penn Center for Health Care Transformation and Innovation and the Center for Health Incentives and Behavioral Economics.

Section Authors

1	Introduction	**Raeesha Altaf**
2	Kickoff	**Raeesha Altaf**
3	Discovery Research	**Raeesha Altaf · Laurel Adams**
4	Understanding and Changing Behavior	**Jeff Ebert**
5	Designing for Scale	**Meg Meador**
6	Prototyping	**Raeesha Altaf · Laurel Adams**
7	Evaluation Plan	**Anish Agarwal · Jeff Ebert**
8	Build and Scale	**Meg Meador**

Design and Illustration

Athena Lee

Editors

Laurel Adams
Arielle Berk
Kit Delgado

First edition, Aug 2026 · Version 1.0
© 2026 The Trustees of the University of Pennsylvania. This toolkit may be reproduced and adapted for non-commercial use with attribution.

Penn Medicine
Nudges in Health Care Toolkit

nudge.unit@pennmedicine.upenn.edu