

# 2

## EDGE EFFECTS: BRIDGING FROM OLD TO NEW

---

Boris Khentov

### 2.1 INTRODUCTION

Ecologists use the term “edge effects” to describe changes that occur at the boundary between two ecosystems. One type of edge effect is increased biodiversity—connecting a new set of dots opens up more opportunities to secure a resource advantage.

The financial services industry, perhaps more than any other industry, is a gnarled mishmash of ecosystems, and no edge is more jagged than the one between the automated, fully digital world and the manual, paper-based world. Increasingly, however, this edge lies just out of sight. Even the stodgiest century-old institutions have invested in web and mobile interfaces, but from a customer experience perspective, that veneer is often skin-deep.

For retail customers coming of age in a world of one-click shopping, same-day delivery, and ride-hailing apps, doing business with financial services incumbents can feel antiquated to the point of disbelief. Transactions we expect to complete with a couple of swipes suddenly grind to a screeching halt.

When you are asked to print, fill out, and mail a form, you might be forgiven for thinking, “Am I being recorded? Is this a prank?” No, this is merely the twentieth century, oddly resilient two decades into the current one, in an industry that

generally attaches little value to nostalgia, at least when it comes to cutting operating costs.

In many corners of retail finance, with no warning, a customer's seemingly routine intent might be whisked into the underworld of legacy back-office workflows, where bits and bytes are inexplicably turned into paper, secured within more paper by a strip of glue, and dispatched across the country (or across the street), only to eventually be revirtualized into disembodied data, roughly five to ten days later.

## 2.2 WOUNDED WORKFLOWS

Retail finance is riddled with such *wounded workflows*, in which jarring disconnects between different generations of technology—one foot in the old and one in the new—are often bridged by menial human labor. At this edge lies both inefficiency and customer frustration, and therefore opportunity.

Continuing to push the line on automating financial services operations may not have the glamour of pursuing the singularity, but it's where many cutting-edge technologies are having the most immediate impact. Some of the more established applications of artificial intelligence (AI) in the private sector include fraud detection, internal compliance, and cyberthreat detection.

However, as in any dynamic environment, not all “edge effects” are productive, and progress can be nonlinear. The seductive power of “the new” is strong, and it pays to remember that the best tool for the job is the one that best solves the problem at hand.

Through my work of leading operations at Betterment, helping to rebuild the investment management industry, I have spent the better part of a decade on the edge of automation. Founded by Jon Stein in 2010, Betterment is a customer-centric fintech money manager that effectively created an entire category that came to be known as *robo-advice*.

In order to grow, an investment management start-up must gather customer assets held by incumbents and has no choice but to accept certain lowest-common-denominator tech realities imposed by the industry. One minor episode from our scaling of Betterment's technical operations illustrates both the challenges and the opportunities of navigating "the edge" as potentially revolutionary new technologies become real options across every domain.

### 2.3 WHERE THERE'S FRICTION, THERE'S IRE

Perhaps the most egregious wounded workflow in investment management is the transfer of an individual's retirement savings between two accounts at unaffiliated institutions, known as a *rollover*. Americans primarily invest for retirement in two types of accounts, IRAs and 401(k)s, which by 2018 held \$9.2 trillion and \$5.3 trillion, respectively.<sup>1</sup>

A technology-enabled newcomer that aims to disrupt this industry, as Betterment did, must not only create a compelling value proposition for potential customers but also be able to logistically process transfers once the assets are in motion. While fully digital transfers are available under certain circumstances, paper checks are inevitable, and with them, the erosion of precision that comes with digital data being expressed in analog form.

There are two primary reasons for this dramatic digital disconnect. The first, not surprisingly, is about incentives. The handful of behemoths that manage custody of the investing public's trillions are not particularly nimble, but mature businesses tend to get *particularly* sluggish when it comes to technology enhancements that make it easier for dissatisfied customers to leave. One need not posit a centralized, inconvenience-loving mastermind at each of these institutions, orchestrating the inertia—work that is likely to decrease revenue in the short term tends to never quite get prioritized.

The second reason is somewhat counterintuitive. Certain financial services processes remain relatively low tech because finance has historically adopted technology eagerly and thoroughly, and even ancient technology can be deemed “good enough” when stability is at a premium. While start-ups have the benefit of starting from scratch, implementing the latest innovations is magnitudes more disruptive to operations when your business is deeply entangled with pervasive legacy technology.

After all, paper checks and the batch systems that print and mail them, and even fiat money itself, were all once a cutting-edge technology, far more efficient at transmitting economic value than, say, taking your chickens to the market and exchanging them for a cow. In a way, this legacy technology remains so entrenched because at some point it was compelling enough to become utterly dominant.

## 2.4 CHECK TECH

The story of how paper checks came to dominate the US financial system is a fascinating one—the key drivers were not only technological but also legal and political. For instance, legislation passed by Congress during the Civil War suppressed privately issued banknotes, which had the effect of encouraging the use of paper checks.<sup>2</sup>

Usage truly exploded in line with post–World War II prosperity, as the number of US checking accounts nearly doubled between 1939 and 1952, kicking off a multidecade operational crisis for the banks that needed to process this ever-swelling tsunami.<sup>3</sup> Some banks had to close their doors at 2:00 p.m. just to get through the backlogs.<sup>4</sup> The surge created a need for a massive pool of workers, tasked with some of the most tedious, mind-numbing clerical work imaginable.

The quest to leverage technology to mitigate the drudgery of check processing began in earnest in the 1950s, when Bank of

America proposed a collaboration with an arm of Stanford University that would pave the way for a number of game-changing developments: a standardized banking industry font that was printed with special magnetic ink, allowing for magnetic ink character recognition (MICR; an early, rudimentary form of optical character recognition<sup>5</sup> [OCR]), and the first automated check-sorting machine employing these technologies.

Subsequent iterations continued to make progress. In 1977, when Cincinnati's regional Federal Reserve bank adopted the latest and greatest sorting machines, processing throughput reached sixty-five thousand checks an hour. The work still required thirty-seven employees, but just a couple of decades earlier, thousands of employees would have been needed.<sup>6</sup>

And yet, much like adding more lanes to a highway has not been shown to reduce traffic, all this added capacity just strengthened incentives to use more of it. While it was easier for consumers to gradually shift over to credit and debit cards for their personal payments, businesses had invested vast sums in check issuing and processing infrastructure, making checks particularly sticky in the context of business-to-business fund transfers—for instance, IRA rollovers.

## **2.5 THE PAPER WEIGHT**

As Betterment's business took off, customers began rolling over their existing retirement savings into their Betterment IRAs. Paper checks issued by other custodians began to stream in, and we had to ensure that each was invested in the correct account, 100 percent of the time, ideally with zero (or at least minimal) human involvement.

We soon discovered that this process was far more challenging than we'd hoped. While checks issued by different institutions follow some general formatting standards (which began with MICR), these do not rise to the level of a "contract" that is sufficiently precise for rules-based software to

process in every instance. The liberties taken by each firm go well beyond font choices.

Our lean operations team was responsible for exponentially more than sorting through this mess, and something had to be done.

## 2.6 TO BUILD OR NOT TO BUILD? THAT IS THE QUESTION

When an operational challenge calls for a technical solution, managers who have the luxury of access to in-house tech talent should ask themselves some version of the “build vs. buy” question.

Real-world business problems are complex, and the answers are often not binary—solutions may get cobbled together and wind up somewhere on the spectrum. But even if the optimal solution were somehow knowable in advance, many companies tend to have an institutional bias, depending on their profile.

To overgeneralize, *traditional* companies (i.e., the kind that refer to their in-house software development teams as IT) tend to be more inclined toward “buying.” When the product calls for substantial customization, that entails some “building,” but this is typically accomplished by “buying” the services of outside consultants.

However, building something out of nothing, precisely to the specifications of the problem, is in every start-up’s DNA. Plus, the requisite talent is generally at hand—when you have a hammer, everything looks like a nail. Therefore, for a business that identifies first and foremost as a technology company, irrespective of vertical (i.e., the kind where IT means ensuring that software engineers, the stars of the show, have ironclad Wi-Fi), the bias might be toward “building” its way out of any challenge.

Like any rapidly growing tech start-up, a fintech business is constantly asking itself, “We could build this ourselves, but *should* we, even if we could do it better?” (Start-ups may be

lacking in resources, but rarely in hubris.) But irrespective of company profile, a good rule of thumb when answering the “build vs. buy” question is to ask yourself two others:

- Is the problem hard? If yes, that’s a vote for “buy.”
- Is the problem important? If yes, that’s a vote for “build.”

There are numerous ways for a problem to be hard. Maybe it calls for an inordinate amount of creativity, and one day some engineer will have her eureka moment and solve the problem. But more likely, particularly in financial services operations, “hard” means messy, fragmented, *boring*, death by a thousand paper cuts (literally in this case).

The problem of generating reliably precise metadata from scanned paper checks is universal across the industry. There are countless companies whose core business is to solve this problem. These are both strong signals for the kind of “hard” you may want to “buy” against.

But few ambitious tech companies can resist the temptation to build against any kind of “hard” if the problem is “important”—another ambiguous descriptor. It’s clearly important enough to address but likely not “important” if the implementation is not core to *your* business strategically (e.g., not a differentiator that will move the needle on growth, or not a potential moat via patent).

For a business that differentiates on customer experience, even a start-up with an itchy “build” trigger finger, back-office automation can be “hard” and usually isn’t “important,” so it tends to tilt toward the “buy” side, provided the cost is reasonable.

As long as your solution works, customers don’t care how elegant it is—not unless you can tell them you’ve eradicated the need for checks entirely. So, we contracted with one of the many vendors whose core business is receiving paper checks and shepherding them back into the digital realm.

## 2.7 FRAGMENTATION ACROSS THE NATION

This so-called lockbox service provided a mailing address and handled physical receipt of rollover checks destined for Betterment. The vendor would open up the envelopes, scan the contents into digital images, and apply some combination of OCR and manual data entry to generate a software-readable metadata attachment. We'd then receive images of the day's checks, along with the metadata, which would include things like the customer's name, the dollar amount of the rollover, and crucially, the destination Betterment account number (since a customer might have several).

There was a “build” component too: internal tooling to ingest this metadata, to allow our operations team to supervise the routing of checks, and to handle exceptions that call for human discretion.

However, the exceptions remained the rule. Because of the maddening fragmentation of the format across different issuers, the metadata from our vendor was nowhere close to 100 percent reliable for full automation, meaning we could not be comfortable with depositing checks with zero human involvement.

Even a customer's name often presented challenges for rules-based automated matching. One custodian might choose to preface the customer's name with “FBO” (for benefit of—an optional legal formality, since retirement accounts are technically trusts). Another might use a customer's nickname, which wouldn't match the full legal name in our database.

Perhaps most frustrating was that each custodian had its own approach for displaying the Betterment destination account number on the check—the single most important data point in determining how to process the rollover. While some have a dedicated field on their checks, many treat it with far less dignity. An exhaustive list of the locations we've seen for account numbers could fill its own (even more boring) chapter:



- Placing it in the “memo” field
- Appending it to the customer’s name in the “payable to” field
- Partially truncating it, resulting in a shorter account number that also happens to exist, which wreaks havoc on automated systems and requires extensive cleanup
- Placing it too close to some other number on the check, causing our lockbox service to mistakenly grab a fragment or the wrong number entirely
- Including it in the “supporting documents” rather than placing it on the check itself, resulting in a blank metadata field

## 2.8 TAXING CONSEQUENCES

At best, a potentially misdirected check would require time-consuming cleanup by our engineering team. At worst, the consequences might be irreversible, depending on how much time had elapsed.

The workings of the US tax code, in all its baroque complexity, call for a particularly high degree of caution. Both IRAs and 401(k) accounts come in two varieties: earnings on which tax is deferred but will be assessed when the funds are withdrawn during retirement (Traditional), and money that is taxed the year it’s earned but can be withdrawn tax-free (Roth).

Many customers have both types of accounts and might initiate two or more distinct transfers around the same time. It’s critical that rollovers between institutions be processed “apples to apples”—inadvertently depositing funds from a Traditional IRA into a Roth IRA could constitute a taxable event for a customer and may go unnoticed for months. Additionally, a rollover from a 401(k) into an IRA, even if the tax profile matches, needs to be reported to the IRS, whereas a transfer between two IRA accounts doesn’t.

None of this important information was printed on various checks in a consistent manner, which called for an approach that is highly intolerant to false-positive matches. Given the operational burden of unwinding an error, you want your automated implementation to err on the side of not making a decision and escalating to a human, rather than risk making the wrong decision. And if you are comfortable automating only about 20 percent of the volume, surely there are more effective uses for your engineering resources.

## 2.9 AUTOMATION AND MORALE

For a lean tech start-up that is used to building to well-defined specifications on top of predictable application programming interfaces, full contact with the messy, physical world of paper checks ignited a seemingly inextinguishable blaze of demoralizing challenges.

The imperative to be deeply attuned to the morale of employees tasked with the most tedious work does not map directly to a dollar amount on your profit and loss statement, yet it's absolutely critical when navigating a multiyear, marathon high-growth scenario.

Our daily volume of rollover checks went from single digits to flirting with breaking a thousand on a busy day. Betterment's operations team gamely grinded under the yoke of this high-stakes yet largely thoughtless work, which seemed out of place in an otherwise dynamic, tech-forward environment.

Automating non-customer-facing functions is rarely given top priority during a start-up's early stages, but it was clear to even the most junior employees that this workflow simply did not "belong" at a financial services business established in the twenty-first century.

On the flip side, automating human workflows, however menial, is a delicate affair. Depending on their mind-set, employees may need varying degrees of reassurance that, by

embracing the process, they will actually enhance their job security—their work will change, but will become more complex and stimulating.

Moreover, managers often have real choices around how to design the scope of the new jobs that emerge after automation. A 2002 MIT study looked at two departments within the same large bank, both of which handled a form of check processing. Both introduced substantial automation improvements around the same time, but the respective managers took different approaches to tasks that were *not* automated.<sup>7</sup>

One department went deeper into specialization, assigning employees to even narrower silos. The other purposefully reshuffled those tasks to create broader-scope roles, designed with the customer in mind, than a single repeatable action.

The latter required a bigger up-front training investment but created higher-paid, more interesting jobs. The findings were consistent with other research, which suggests that such integrative job designs are more commonplace where management attaches importance to increasing the well-being of employees and where customer service is a critical goal.

We were fortunate in being able to attract ambitious, talented people into junior roles and knew they'd have no qualms with solving more customer problems—and solving them faster and with more leverage. The entire team could band together and focus on automating as much of this work as possible.

## **2.10 MORE DATA, MORE PROBLEMS**

Spoiled by the absence of legacy data sets of inconsistent integrity and now forced to absorb the externalities imposed on us by legacy technology outside our control, we stumbled into a classic record linkage (RL) problem (also called “data matching”), studied extensively beginning in the 1950s.<sup>8</sup>

The need to create a link between different representations of the same object across multiple data sources (using one or

more “common identifiers”) is a typical RL workflow. In our case, the unique object was the customer’s rollover transaction. This was represented in three distinct data sources, with a number of “common identifiers” available, but none was reliably sufficient on its own, without reference to others.

Betterment’s systems had a representation of this transaction, as did the other custodians. By issuing a paper check, the custodian created a third representation—a particularly troublesome one, when it came to reliably extracting the common identifiers.

While RL problems are as ancient as the IBM mainframes that “nobody ever got fired for buying,” machine learning has opened up a new class of solutions for automating RL decision-making in the twenty-first century. Rather than linking on one or two unambiguous common identifiers, overall certainty is achieved by weighting the certainty signals across a potentially wide set of common identifiers. This is called *probabilistic record linkage* or *fuzzy matching*—a seemingly ideal term for the chaotic, unpredictable world of typos, smudges, and formatting shenanigans.

## **2.11 TO LABEL IS HUMAN, TO STORE, DIVINE**

While tech start-ups are perpetually strapped for resources, they offset this by conferring a number of advantages, including the much-touted ability to make decisions rapidly. A more understated advantage is the physical proximity of technical experts to those in nontechnical roles, which leads to spontaneous interactions and serendipitous efficiency gains.

That’s precisely what happened when Sam Swift, then Betterment’s head of data science, looked over the shoulder of a junior operations associate and was intrigued by what he saw. The associate was using our internal tool, which loaded the metadata that the vendor pulled from the checks, and made a recommendation for which Betterment account each check should be deposited into. The associate was flipping through

one scanned image after another, confirming each decision, occasionally pausing to peer at one for longer than the others.

After a few questions, Sam identified a classic machine learning problem, which consists of the following conditions:

1. Well-defined set of inputs (name, dollar amount, source account type, destination account number)
2. Easy to define “success” (check deposited into the right account)
3. Large data store of “human-labeled” outcomes

The third condition is particularly important, and routinely underappreciated.

When considering pricey investments in AI, business leaders readily recognize that their primary data needs to be in digital form and readily accessible. What’s less intuitive without a technical background is the need to have been creating a massive record of past decisions *made by humans* based on that primary data, particularly for “matching” problems.

Many of us first experienced the power of a massive human-labeled data set during that disorienting moment when Facebook’s “tag your friends” feature suddenly began making eerily precise “recommendations” as to the identity of every person in a photo, trained on untold millions of our own, manual selections.

In many instances, however, an adequate human-labeled data set simply does not exist. This winds up being one of the primary blockers to implementing machine learning solutions, even at large companies with ample budgets, whose leadership is sold on the investment.

## 2.12 A BOT IS BORN

In our case, by the time Sam’s interest was piqued, humans at Betterment had already looked to the same fields to process hundreds of thousands of checks. A machine learning model would

train on years' worth of human decisions and "learn" what constitutes a "match" between the occasionally messy metadata pulled from a check and the record of the transaction already in our system.

Mystical terms such as "gradient boosting," "random forests," and "stacked ensembles" began to swirl in a corner of the Betterment office unaccustomed to such wizardry. Incantations of the AI age, they carried hope for lifting the curse of unstructured text fields and fat-fingered typos, and all the frustration they brought.

Thus was conceived the Betterment Rollover Bot, powered by AI. The plan was to run the bot in parallel with the human-supervised process and observe it, before gradually shifting any actual decision-making over to the machine.

In the summer of 2017, the Rollover Bot became sentient, purely in "test mode." Something about this prosaic, unassuming application of machine learning, deep under the hood, where no customer would ever experience it, felt symbolic. While we were patiently waiting for AI to drive our cars and generally "change life as we know it," was it already changing the world in less flashy ways, one healed "wounded workflow" at a time?

A couple of months later, I presented this use case at a fintech industry committee meeting, as an illustration of real-world applications of AI in start-up operations. I ended my talk on a high note, hopeful for additional modest breakthroughs.

### **2.13 KEEP IT SIMPLE, SOMETIMES**

Fast-forward two years, to 2019. David Shrier, who sat on the same committee and had seen the presentation, reached out and asked if Betterment would be interested in contributing a chapter to this book. He suggested that the Rollover Bot story would be a good fit.

My first thought was, “Well, this is awkward,” because as much as we appreciated the opportunity, the Rollover Bot had its last test run a few months earlier, having never made the jump to production. Retelling the story today, while leaving that part out, didn’t feel right. However, as I mulled it over, it occurred to me that the *whole* story might be an even better fit, so here’s what happened next.

Back in 2017, when the Rollover Bot began its daily test runs, our operations team eagerly looked on, anticipating an imminent, joyful handoff of its most thankless work. Yet, something wasn’t clicking the way we expected.

The “match rate” was initially promising, but the Rollover Bot never reached a level of confidence in which it could relieve the humans of even half the volume. It’s almost as though something was missing—something that humans intuitively know to reference when resolving a variety of edge cases, which wasn’t available on the check alone.

A couple of months after the Rollover Bot began its lackluster life, AJ Kramer, Betterment’s then-director of operations, had an idea. Comparatively speaking, it was about as high tech as a bag of rocks, but it was in production within a couple of months, and it fully automated nearly 80 percent of Betterment’s rollover check volume.

AJ’s idea was as simple as it was effective. Rollovers into Betterment are initiated inside our app. When customers express their intent to transfer an IRA to us, they also provide some basic information as part of the flow, which then triggers instructions that are sent to the customer’s current custodian. For each rollover *intent* created in the app (customers can initiate several, from and into multiple accounts) we would generate a new, unique alphanumeric ID that began with the three-letter prefix “BMT.”

We instructed the transferring institution to put this “rollover intent ID” in the memo field of the paper check (see figure 2.1). We then instructed our lockbox service to look for

WARNING: CHECK HAS COLORED BACKGROUND WITH COLORED, VISIBLE FIBERS AND A PINK SPOT IN BOTTOM RIGHT HAND CORNER THAT REACTS TO HEAT. DO NOT CASH WITHOUT THESE FEATURES.



Principal Trust Company  
Des Moines, IA 50392-0001

70-2392/719

Date  
NOT VALID AFTER 180 DAYS

Amount \$ \*\*\*\*\*16,344.73\*



Pay To

ACCT# BMTW950W80  
Betterment  
P.O. Box 203945  
Dallas TX 75320

14

SIXTEEN THOUSAND THREE HUNDRED FORTY FOUR AND 73/100-----

THE NORTHERN TRUST COMPANY, CHICAGO, IL 0710  
PAYABLE THROUGH NORTHERN TRUST BANK/DUPAGE, OAK BROOK, IL



By

Authorized Signature(s)

FIGURE 2.1

Rollover intent ID on a paper check.



this ID and to include it as an additional field in the metadata passed on to us. We could then look *solely* at this ID and match the check to our record of the rollover intent (which includes the destination Betterment account number). If this ID was present in the metadata, we could be confident of a perfect match.

## 2.14 THE MISSING LINK

In RL parlance, AJ had created a new common identifier, which for some reason was dramatically more effective than the existing ones at linking records, on its own, without reference to the others. What advantage did it have over the Betterment destination account number, also a unique ID, which wasn't getting the job done? It turns out there were three factors that, in aggregate, materially reduced error rates:

1. Check issuers know what a destination account number is, and each has its own process for it, driving fragmentation. **Injecting an additional, “unknown” string into the process actually resulted in far more predictable placement on the check.** We asked them to put it in the memo field, and much of the time they did, in the absence of a conflicting internal protocol.
2. Because it was alphanumeric, with a unique BMT prefix, **the new identifier was purposefully designed to “pop” better among a bunch of numbers.** Through whatever mix of OCR and manual data entry our lockbox service employed, it made far fewer mistakes parsing it out.
3. Subtly, but perhaps most importantly, **the new identifier had a one-to-one correlation with the record we were trying to link to** the record of a specific transaction, not the record for the destination account. The account is a good proxy, but it can receive multiple rollovers over the same time frame, from different account types (with

distinct tax implications). A unique identifier, created during the session when the customer expresses the rollover intent, would point to all the data that was available when the transaction record was created (but no superfluous data, which could create ambiguity).

When researchers tag animals before releasing them back into the wild, they are effectively creating *state awareness*, preparing the system to derive meaning when reencountering the same specimen, linking the new impression with the existing record. Similarly, we didn't need more sophisticated pattern recognition across a complex set of signals—just a more reliable, solitary tag to link us back to something we already knew.

As for the Rollover Bot, it continued running in test mode for another year or so. It cost us nothing to do this, and we were either too busy to put it to sleep or morbidly fascinated with observing its now purposeless existence, or a bit of both. But in our minds, it will forever symbolize the somewhat corny but undeniably wise principle of “fall in love with the problem, not the solution.”

## 2.15 WHAT'S YOUR PROBLEM?

So, what happened, exactly? We encountered an easier category of the RL problem masquerading as a harder one, which we might have unmasked sooner, had it not been for our eagerness to deploy the “cool” technology.

Our initial assessment of the problem was informed by the fact that routing a check often requires human judgment because the current process cannot parse checks with accuracy and consistency. Under this framing, reaching for the best technology available that attempts to replicate human “intelligence” seems reasonable. However, there were two implicit assumptions:

- The check contains all the information necessary to make a decision.

- The visual profile of another firm's check is a fixed constraint.

Thanks to his domain expertise, AJ couldn't help but challenge those assumptions.

- Observing his team's wounded workflow, AJ noted that, even after closely studying a "problem check," a human (that paragon of judgment we sought to replicate) would sometimes reach for other data elsewhere in the system to resolve an ambiguity. Something was missing.
- He was also familiar with the check-issuing processes of the big custodians and recognized that we have a limited way to indirectly influence what goes on a check issued by a third party.

He restated his assessment of the problem as: routing a check often requires human judgment because the current process cannot parse checks with accuracy and consistency.

*However:*

- Even when parsed perfectly, a check may not contain all the information required to make a decision (though that info is available elsewhere in our system), and
- The visual profile of another firm's check is not entirely immutable.

In this case, the heavy lifting was in assessing the true parameters of the problem. An accurate assessment, made possible through creativity and domain expertise, made the problem "easier."

Though he wouldn't have known the lingo, AJ had turned a "fuzzy matching" RL problem, which calls for fancy AI models to pattern match across multiple identifiers, into the simplest, rules-based RL problem, matching with high confidence on a single common identifier (which he first had to invent).

It was an elegant work-around, which is another way of saying "a better solution."

## 2.16 THE RIGHT STUFF FOR THE RIGHT KIND OF “NEW”

What made this improvement (and many of our other efficiency gains) possible wasn't the latest and greatest tech. Incremental breakthroughs are often about new connections between largely familiar points of reference, which flourish when there's an opportunity to rethink problems from scratch.

Generating a unique “intent identifier” associated with a customer-driven user experience flow, solely for the purpose of solving an RL problem, thereby automating a back-office workflow, is domain-specific to the point of tears and not inherently instructive of anything.

What's notable is the set of circumstances under which such an idea would actually be implemented. In those rare moments when it becomes viable to rebuild something massive (like a money manager) from scratch, all kinds of structural assumptions begin to fall, and things just start working better.

The primary challenge facing incumbent financial institutions today isn't technology per se; it's reorienting toward a holistic customer-centric experience. These giants are throwing huge sums at creating modern interfaces, but the teams responsible for these experiences, no matter how talented, simply do not have the scope and agency to control the *entire* experience, which is also dependent on legacy systems that have existed for many decades. Matching the experiences offered by start-ups requires the will and ability to meaningfully bridge these two worlds, which can be more of an organizational question than a technological one.

## 2.17 MEANS, NOT ENDS

No corner of human activity is devoid of parables—not even a humble fintech back-office operation. Those of us who have

devoted our careers to tech should take great care to remember that technology, in and of itself, isn't the point.

In fact, one of my favorite pastimes when working with technical teams is to ask if we've explored every option to "make the problem go away" without resorting to the tools we're familiar with (or the tools we'd *like* to be more familiar with).

For what it's worth, machine learning has found better (more suitable) applications at Betterment and will surely proliferate to other use cases as well. But I'll always remember the Rollover Bot as a symbol of how hype and desire for "the new" carries with it the risk of distorting judgment and losing sight of what actually matters.

## NOTES

1. G. Iacuri, "IRA Assets Will Almost Double Those in 401(k) Plans over Next Five Years," *InvestmentNews*, June 25, 2018, <https://www.investmentnews.com/ira-assets-will-almost-double-those-in-401k-plans-over-next-five-years-74724#:~:text=Assets%20in%20individual%20retirement%20accounts,according%20to%20an%20industry%20analyst>.
2. S. Quinn and W. Roberds, "The Evolution of the Check as a Means of Payment: A Historical Survey," Federal Reserve Bank of Atlanta *Economic Review* 93, no. 4 (2008), [https://www.frbatlanta.org/-/media/documents/research/publications/economic-review/2008/vol93no4\\_quinn\\_roberds.pdf](https://www.frbatlanta.org/-/media/documents/research/publications/economic-review/2008/vol93no4_quinn_roberds.pdf).
3. D. Scharfenberg, "Paper Checks Forever?," *Boston Globe*, April 21, 2018, <https://www.bostonglobe.com/ideas/2018/04/21/paper-checks-forever/B9KYFwuXGOieb0A35CAEdN/story.html>.
4. E. Smith, "Scanning and Sorting," *Tedium*, February 17, 2021, <https://tedium.co/2021/02/17/check-processing-history-automation/>.
5. E. Smith, "Seek and Spell," *Tedium*, March 22, 2017, <https://tedium.co/2017/03/22/ocr-typography-optical-character-recognition-history/>.

6. Staff, "At Cincinnati Fed, \$200 Million Is a Normal Night," *Cincinnati Enquirer*, September 11, 1977.
7. D. Autor, F. Levy, and R. J. Murnane, "Upstairs, Downstairs: Computers and Skills on Two Floors of a Large Bank," *Industrial and Labor Relations Review* 55, no. 3 (2002): 432–447.
8. H. L. Dunn, "Record Linkage," *American Journal of Public Health* 36, no. 12 (1946): 1412–1416.