Kenneth O. Stanley · Joel Lehman

Why Greatness Cannot Be Planned

The Myth of the Objective

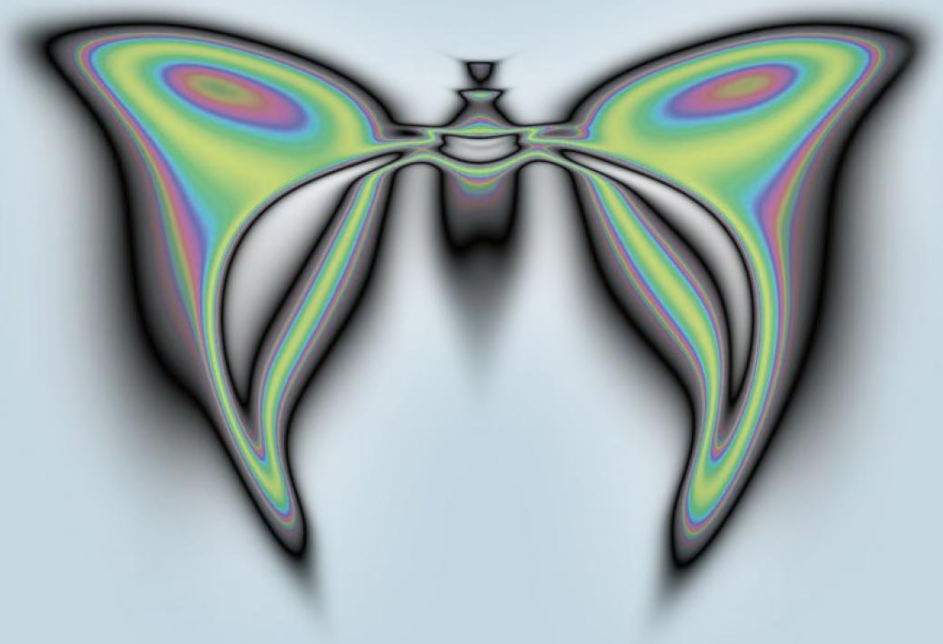

Springer

Kenneth O. Stanley • Joel Lehman

Why Greatness Cannot Be Planned

The Myth of the Objective

Springer

Kenneth O. Stanley
Department of EECS
Computer Science Division
University of Central Florida
Orlando, FL, USA

Joel Lehman
Department of Computer Sciences
The University of Texas at Austin
Austin, TX, USA

The image on the cover has been used with permission from the University of Central Florida Research Foundation, Inc. It was created using the www.picbreeder.org website of UCFRF.

ISBN 978-3-319-15523-4 ISBN 978-3-319-15524-1 (eBook)
DOI 10.1007/978-3-319-15524-1

Library of Congress Control Number: 2015932228

Springer Cham Heidelberg New York Dordrecht London

Printed on acid-free paper

Springer International Publishing AG Switzerland is part of Springer Science+Business Media (www.springer.com)

Contents

Chapter 1
Questioning Objectives

> *Now, in his heart, Ahab had some glimpse of this, namely: all my means are sane, my motive and my object mad.*
>
> Herman Melville, *Moby Dick*

Imagine waking up without thinking about what you need to do today. Have you ever done that? Suppose then you go to work, but instead of holding the usual meetings and poring over benchmarks and milestones, your boss instead tells you to *do whatever you feel is most interesting*. What would you do? When you later check the news online, there is no talk of national testing standards or missed objectives for the economy. Strangely enough, teachers somehow still teach and money somehow still changes hands anyway. Maybe you post a profile on a dating website, but you leave blank all the answers on what you look for in a partner. Today you aren't looking for anything in particular, but that still doesn't stop you from looking. While you probably won't experience a day like this one anytime soon—a day without objectives—what would life be like if you could? Would it be confusing, difficult to navigate, too open-ended—or, quite the opposite, a lot *better*?

It's interesting that we rarely talk about the dominance of objectives in our culture even though they impact us from the very beginning of life. It starts when we're barely more than a toddler. That momentous first day that we enter kindergarten is the gateway to an endless cycle of assessment that will track us deep into adulthood. And all that assessment has a purpose—to measure our progress towards specific *objectives* set for us by society or by ourselves, such as mastering a subject and obtaining a job. The reality is that objectives lurk in the background from the earliest moments. But they only start there. Over the years they keep gathering steam, eventually enveloping practically everything we do.

If you want evidence for that, all you need to do is stop in at your local bookstore and take a look at the magazine rack. There you will be reminded that you might want to *change jobs*, *lose 15 pounds*, *start a company*, *find a date*, *get promoted*, *change your look*, *make a million dollars*, *buy a house*, *sell a house*, or even *complete a video game*. In fact, almost anything worth doing is expressed as an objective. Now, we're not suggesting by any means that all of these objectives are a waste of time. Of course, many are admirable. Others may be more frivolous, but whatever you may think of one objective or another, we rarely question the value of

K.O. Stanley, J. Lehman, *Why Greatness Cannot Be Planned*,
DOI 10.1007/978-3-319-15524-1_1

framing all our pursuits with objectives. Can you imagine life with fewer objectives, or with no objectives at all? Would it lead to any good? Regardless of how you'd answer these questions, at the least they highlight how deeply our culture has come to revere objectives.

And it's not only about individual pursuits. While it's true that children are graded for their progress towards mastering a subject, the schools themselves are *also* graded. But in the case of the schools, their objective is to produce high student test scores. Even nations set objectives, like low crime, low unemployment, and low carbon emissions. A lot of effort and resources are spent measuring progress towards these objectives and others like them. There's an assumption behind these pursuits that isn't often stated but that few would think to question: We assume that any worthy social accomplishment is best achieved by *first* setting it as an objective and *then* pursuing it together with conviction. It makes you wonder, is there such a thing as accomplishment without objectives?

If you take a look at most professions, the answer would appear to be no. Take engineers for example. They typically set objectives through rigorous specifications. Then they continually measure how their prototypes compare to these specifications. Inventors are similar, conceiving an invention and then pursuing it as an objective. In the same spirit, scientists must come up with clear objectives to secure funding for their projects, which are then judged by how likely they are to achieve their objective. The list goes on. Investors set earnings objectives and corporations set profit objectives. Even artists and designers often mentally conceive a design and set its realization as their objective.

The weight of objectives on our thinking even impacts the way we talk about topics like animals in nature. After all, whenever evolution is discussed, we view animals through the lens of *survival and reproduction*, evolution's assumed objective. Even many *algorithms*, or programs that run inside computers, are designed to work towards some specific objective, such as finding the best search result or playing a better game of chess. In fact, these kinds of algorithms are quite common in artificial intelligence and machine learning. So it may come as no surprise that the term *objective function* is now practically a household name in such fields.

Maybe all this objective mania makes sense. At some level, we must believe it does to have allowed it to dominate our lives so completely. Or maybe it's something else—maybe we've become so used to objectives defining everything we do that we've forgotten that their value can even be questioned. Either way, there's a certain appeal to this kind of routine. The idea that all our pursuits can be distilled into neatly-defined objectives and then almost mechanically pursued offers a kind of comfort against the harsh unpredictability of life. There's something reassuring about the clockwork dependability of a world driven by tidy milestones laid out reliably from the starting line.

Though often unspoken, a common assumption is that the very act of setting an objective creates possibility. The very fact that *you put your mind to it* is what makes it possible. And once you create the possibility, it's only a matter of dedication and perseverance before you succeed. This can-do philosophy reflects how deeply

optimistic we are about objectives in our culture. All of us are taught that hard work and dedication pay off—if you have a clear objective.

Even so, perhaps you've felt qualms about this kind of thinking from time to time. It might sound good on the face of it, but what it *leads to* isn't always so comforting—legions of measurements, assessments, metrics—woven into every aspect of our lives. It's like we've become slaves to our objectives, toiling away towards impossible perfection. Objectives might sometimes provide meaning or direction, but they also limit our freedom and become straitjackets around our desire to *explore*. After all, when everything we do is measured against its contribution to achieving one objective or another, it robs us of the chance for playful discovery. So objectives do come with a cost. Considering that this cost is rarely discussed in any detail, maybe it's a good idea to look a little harder at what we're really giving up in exchange for such objective optimism.

But before we do, it's important that you know that we aren't pessimists. It may sound like this book is going in a cynical direction, but that isn't really true. In fact, we believe that human achievement has no limits. It's just that we're going to highlight a different path to achievement, without the need for objectives. There's a lot our culture has sacrificed in the name of objectives, and we're going to take it back. They've stolen our freedom to explore creatively and blocked us from serendipitous discovery. They ignore the value of following a path for its own uniqueness, rather than for where it may lead. The chapters ahead will show that great discoveries are lurking just beyond our fingertips, if only we can *let go* of the security blanket of the objective. Sometimes, the best way to change the world is to stop *trying* to change it—perhaps you've noticed that your best ideas are often those you were not seeking. We'll come back to this paradox later, but first let's think a little more about the way most people approach achievement today.

* * *

It usually starts with deciding on an objective. In a lot of professions, the first question you'll hear if you propose a new project is, "What is the objective?" People often say that a particular pursuit is not "well-defined" enough unless you can tie it to an objective. That's usually the only way to prove your idea is worth considering. For example, even if you have a strong hunch that combining two chemicals will lead to an interesting reaction, hardly any scientist will take you seriously unless you can *define* what that interesting reaction would be. Only then could you say that you have a clear objective, and a legitimate pursuit is underway.

Sometimes a word other than "objective" might be used, but usually it plays a similar role. For example, scientists often demand *hypotheses* from each other because they don't want to fund research only because it sounds "interesting." "What is the hypothesis?" they'll insist. It's similar to asking what is the objective. Without a hypothesis, an experiment is reduced to *mere speculation*, little more than child's play. The hypothesis, like the objective of any project you might want to pursue, certifies that it's worth pursuing. Even if it might not work out, it still provides a clear outcome that will allow people to judge your success or failure later.

This attitude doesn't only apply to business or science. It's personal too. If you join a dating site, you're supposed to know what you want and what you like, so that you can describe what you're "looking for." If you're starting college, you're supposed to figure out your major so that you have a clear body of knowledge to pursue. And as any high school student realizes who is applying to college, even hobbies are supposed to have *some* kind of purpose. "Staring at the wall" will not meet much respect on a college application, even if it's when you have your best ideas.

Whether you're an executive, a scientist, a student, or even just a single looking for a date, once your objective is defined usually the next step is to put all your energy into achieving it. In other words, we pull out all the stops in the pursuit of our goal. But there's actually one other ingredient in pursuing an objective that is not as obvious but almost always present: Usually progress towards the objective is somehow *measured*. This is where all the measurements and metrics of our culture come into play. The purpose of all the measuring is to help us figure out whether we're heading in the right direction. If things aren't going well, it gives us the chance then to switch directions. For example, in school the objective is to master a subject. Grades are a measure that tell us whether we're making progress in that direction. So if your grades in school are going down, you may need to change your approach to studying.

There's a useful word that comes from optimization theory that some scientists use to describe this idea of using a measure to help decide what to do next—they call the measure a *gradient*. It's basically a clue to which direction is right and which is wrong. We all quickly become very familiar with following such gradients, to the point that it almost becomes unconscious by the time we're adults. The "hot and cold" children's game is a perfect introduction to the idea: The young contender searches for a hidden treasure known only to the other players, who provide clues by saying "hotter" or "colder." The idea is so simple that children naturally follow the "gradient" of increasing temperature with little need for explanation (and definitely no need to learn anything about optimization theory). As long as the gradient is going up, chances are that success is coming closer. In a sense we never stop playing the hot and cold game. The process of setting an objective, attempting to achieve it, and measuring progress along the way has become the primary route to achievement in our culture.

Obviously we're going to raise a lot of questions about the benefits of objectives, but one important point is that we're mainly focused on *ambitious* objectives—those whose achievement is anything but certain. One of the reasons that objectives aren't often questioned is that they work perfectly well for more modest pursuits. If a manufacturer decided to increase efficiency by 5 %, no one would be shocked if it succeeded. A software company upgrading its product from version 2.0 to version 3.0 is similarly likely to succeed, as happens all the time. Everyday successes like these mislead us into believing that setting objectives works well for almost *everything*. But as objectives become more ambitious, reaching them becomes less promising—and that's where the argument becomes most interesting.

Some objectives are anything but certain. Medical researchers have yet to cure cancer. And it's hardly clear that computer scientists will succeed in creating a convincingly-human artificial intelligence any time soon. It would be nice to have a limitless source of energy with no risk or environmental impact whatsoever, but who knows when that will happen. Even fully immersive holographic television, which you must admit would be fun though perhaps less noble, is nowhere in sight. Maybe there is even a beautiful new genre of music waiting to be found, so enchanting that all of humanity would succumb to its spell, if only the right artist can find it. Or, if you want to get *really* ambitious, how about time travel or teleportation? You might even have some big ones of your own, like making a billion dollars. Of course some ambitious objectives might turn out impossible. But *some* are possible, and if we could even come close to achieving them, the world would surely change for the better.

The question then is, *how will we ever get there?* How can we achieve not just improving our golf swing, but our *dreams*? These are the kinds of questions where objectives offer less assurance. After all, even if time travel is theoretically possible, the best use of national resources at this moment is surely not to reroute trillions of dollars into an all-out effort to build a time machine. But why not? Is setting the objective not the first step towards success? Or would pursuing time travel be our Moby Dick, distracting us from what really matters in life?

* * *

Why is it so hard to achieve ambitious objectives? To get a handle on the problem, think about all the things that are *possible*. As an example, consider all possible *images*. Everything you've ever seen falls into this class, but so do many things you've never seen and will never see. Now consider this: A minuscule proportion of all possible images are great masterpieces like Leonardo da Vinci's *Mona Lisa* or Vincent Van Gogh's *The Starry Night*. These are objectives that are hard to achieve. A larger proportion of possible images are recognizable but less inspiring than the masterpieces, like pictures of acquaintances and everyday objects. Of course, the vast majority of possible images are of no interest whatsoever—like television screens tuned to the wrong channel, just random static. One exciting aspect of this vast set of all possible images is that some of the great masterpieces *do not yet exist* because no one has painted or drawn them yet. Put another way, these future masterpieces are not yet *discovered*.

It's useful to think of achievement as a process of discovery. We can think of painting a masterpiece as essentially discovering it within the set of all possible images. It's as if we are *searching* through all the possibilities for the one we want, which we call our objective. Of course, we're not talking about search in the same casual sense in which you might search for a missing sock in the laundry machine. This type of search is more elevated, the kind an artist performs when exploring her creative whims. But the point is that the familiar concept of search can actually make sense of more lofty pursuits like art, science, or technology. All of these pursuits can be viewed as searches for something of value. It could be new art, theories, or inventions. Or, at a more personal level, it might be the search for the right career.

Whatever you're searching for, in the end it's not so different from any other human process of discovery. Out of many possibilities, we want to find the one that's right for us.

So we can think of creativity as a kind of search. But the analogy doesn't have to stop there. If we're searching for our objective, then we must be searching *through* something. We can call that something the search *space*—the set of all possible things. Now try to picture this space—it's as if different possibilities appear at different locations in a big room. Imagine this giant room in which every image conceivable is hovering in the air in one location or another, trillions upon trillions of images shimmering in the darkness, spanning wall to wall and floor to ceiling. This cavernous room is the space of all possible images. Now imagine walking through this space. The images within it would have a certain *organization*. Near one corner are all kinds of faces, and near another are starry nights (and somewhere among them Van Gogh's masterpiece). But because most images are just television noise, most of the room is filled with endless variations of meaningless gibberish. The good stuff is relatively few and far between.

The nice thing about thinking of discovery in terms of this big room is that we can think of the process of creation as a process of searching through the space of the room. As you can imagine, the kind of image you are most likely to paint depends on what parts of the room you've already visited. If you've never seen watercolor, you would be unlikely to suddenly invent it yourself. In a sense, civilization has been exploring this room since the dawn of time. As we explore more and more of it, together we become more aware of what is possible to create. And the more you've explored the room yourself, the more you understand where you might be able to go next. In this way, artists are searching the great room of all possible images for something special or something beautiful when they create art. The more they explore the room, the more possibilities open up.

Let's pretend you wanted to paint a beautiful landscape—so that's your objective. If you're experienced in landscape painting, it means that you've visited the part of the room teeming with images of landscapes. From that location, you can branch off to new areas full of landscapes that are still unimagined. But if you're unfamiliar with landscape painting, unfortunately, you're unlikely to create a masterpiece landscape even if that's your objective. In a sense, the places we've visited, whether in our lives or just in our minds, are *stepping stones* to new ideas.

This way of thinking doesn't only work for paintings. We can imagine a room filled with anything. For example, it could be a giant room full of inventions. As with images, the great room of all possible inventions is mostly filled with inelegant doodads, like unreliable machines that occasionally can move a ball from one location to another and pointlessly ring a bell. But if you look hard enough, some areas of the room contain simple yet useful ideas—bags, wheels, wheelbarrows, spears. Rarer still are more sophisticated inventions—in one corner are all manner of cars and in another corner, computers. Somewhere in the room are fanciful machines unlike any we've seen before, still waiting to be found.

Let's visit the pocket of space filled with computers. If you spend a lot of time in this area, an interesting thing begins to happen: You start to understand the shape of

the space, how one computer leads to another, like stepping stones along a winding path. Wander here long enough, and you might even start to see where the interesting possibilities may lie.

But why do we have to wander? Why can't we just go directly to the location of the most powerful machine of all? The reason is that the only clue to the future is found in the past. The first computer, the ENIAC, was introduced in 1946 and could execute 5,000 instructions per second [1]. While this might sound like a lot, a modern desktop computer executes over *ten billion* instructions per second [2]. In other words, ENIAC was a slug, more than two million times slower than today's computers.

You might wonder, why didn't the designers simply set their objective to build a fast computer in 1946? After all, we now know it's possible. Why start so slow? But the world always works this way. Before anyone had explored the part of the great room filled with computers, no one could know what was even possible there. You have to visit before you can depart. In short, the stepping stones to faster computers were not yet known in 1946 because they had not yet been *discovered*, just as we do not yet know the stepping stones that may lead to computers two million times faster than today's. Stepping stones are portals to the next level of possibility. Before we get there, we have to find the stepping stones.

The moral of course is that objectives have to be realistic for us to have the hope of reaching them. The idea of a large imaginary "room" whose space is being searched is a metaphor that helps to see why this principle holds. It's no surprise then that computer scientists even have a term, *search space*, that refers to this very concept—it's the general idea that creation and discovery happen within a space of possibilities that contains stepping stones leading from one discovery to another. The big question when setting ambitious objectives is always whether we can find a path through this search space from where we currently are to where we want to be in the future—the path to the objective.

Sometimes, figuring out the path from here to the objective isn't exactly a major challenge. For example, if your objective is to eat, you probably won't need to explore very much to figure out that you need to visit the refrigerator before making a sandwich. But the paths to many objectives are of course a lot more difficult to guess. If you're a freshman in college, there may be a path to making a million dollars before you reach 30, but what you should do first is probably not so obvious. In other words, you'll have to *search* for the right stepping stones, and if you're lucky and clever enough, you might discover the ones that lead to the objective. There may be a number of stepping stones you have cross, and many of them are likely challenging to figure out.

* * *

Now we can finally tell you what this book is really about. It's not just about achievement or how to succeed. The deeper message here is about a surprising paradox. We'll have a lot more to say on it in later chapters, but you can get the gist of it just from what we've said so far: Objectives are well and good when they are sufficiently modest, but things get a lot more complicated when they're

more ambitious. In fact, objectives actually become *obstacles* towards more exciting achievements, like those involving discovery, creativity, invention, or innovation—or even achieving true happiness. In other words (and here is the paradox), the greatest achievements become *less likely* when they are made objectives. Not only that, but this paradox leads to a very strange conclusion—if the paradox is really true then the best way to achieve greatness, the truest path to "blue sky" discovery or to fulfill boundless ambition, is to have *no objective at all.*

Why would this be? The key problem is that the stepping stones that lead to ambitious objectives tend to be pretty strange. That is, they probably aren't what you would predict if you were thinking only of your objective. Put another way, if you were walking through one of these great rooms of all possible things, how the items in the room are arranged would be unpredictable and confusing. History is full of examples of this vexing problem. For example, the very first computer was built with vacuum tubes, which are devices that channel electric current through a vacuum. However, here's the strange part: The history of vacuum tubes has nothing to do with computers. People like Thomas Edison who were originally interested in vacuum tubes were investigating electricity, not computing. Later, in 1904, physicist John Ambrose Fleming refined the technology to detect radio waves, still with no inkling of building a computer. It was only decades later that scientists first realized that vacuum tubes could help build computers, when the ENIAC was finally invented [3].

So even though vacuum tubes are a key stepping stone on the road to computers, few if any could see it coming. In fact, if you were alive in 1750 with the objective of building some kind of computer, you'd never think of inventing a vacuum tube first. Even *after* vacuum tubes were first discovered, no one would realize their application to computation for over 100 years. The problem is that *the stepping stone does not resemble the final product.* Vacuum tubes on their own just don't make people think about computers. But strangely enough, as history would have it vacuum tubes are right next to computers in the great room of all possible inventions—once you've got vacuum tubes you're very close to having computers, if only you could see the connection. The problem is, who would think of that in advance? The arrangement, or *structure*, of this search space is completely unpredictable.

Unfortunately, this kind of unpredictability is the rule rather than the exception in almost any situation with an ambitious objective. The first engine was not invented with airplanes in mind, but of course the Wright brothers needed an engine to build a flying machine. Microwave technology was not first invented for ovens, but rather was part of magnetron power tubes that drove radars. Only when Percy Spencer first noticed the magnetron melt a chocolate bar in his pocket in 1946 did it become clear that microwaves are stepping stones to ovens [4].

These stories of delayed revelations and serendipitous discovery expose the danger of objectives: If your objective was to invent a microwave oven, you would not be working on radars. If you wanted to build a flying machine (as countless failed inventors did over the years), you wouldn't spend the next few decades instead trying to invent an engine. If you were like Charles Babbage in the 1820s [5] and wanted

to build a computer, you wouldn't dedicate the rest of your life to refining vacuum tube technology. But in all these cases, what you would never do is exactly what you should have done. The paradox is that the key stepping stones were perfected only by people *without* the ultimate objective of building microwaves, airplanes, or computers. The structure of the search space—the great room of all possible things—is just plain weird. It's so bad that the objective can actually *distract* you from its own stepping stones! If you think too much about computers you'll never think of vacuum tubes. The problem is that ambitious objectives are often *deceptive*. They dangle a false promise of achievement if we pursue them purposefully. But strangely in the end we often must give them up ever to have the chance of reaching them.

This paradox isn't just about historical events. It applies just as much today as ever—to everything from society's greatest challenges to your own personal ambitions. Chances are that if we plan a path based on our objective, then it will miss the stepping stones. This insight raises some troubling questions that relate to the tricky nature of stepping stones: Do increasing test scores really lead to subject mastery? Is the key to artificial intelligence really related to intelligence? Does taking a job with a higher salary really bring you closer to being a millionaire? Is cancer going to be cured by an insight from someone who is not a cancer researcher? Are improvements in television technology bringing us any closer to holographic television?

It often turns out that the *measure* of success—which tells us whether we are moving in the right direction—is deceptive because it's blind to the true stepping stones that must be crossed. So it makes sense to question many of our efforts on this basis. But actually the implications are even more grave than just questioning particular pursuits and their objectives. At a deeper level, we might ask why we think ambitious pursuits should be driven by objectives *at all*.

* * *

It's not hard to find great ideas that were never objectives for anyone, at least not until almost the moment they were discovered. Rock and roll music took inspiration from jazz, blues, gospel, and country music. In a way, those genres acted as stepping stones to rock and roll. But no one was *trying* to discover rock and roll because no one knew it even existed as a possibility. Jazz musicians weren't trying to influence the birth of rock any more than ragtime composers consciously tried to shape jazz. Even so, strangely enough, ragtime *did* shape jazz and jazz did shape rock and roll.

Count Basie, who was a respected name in jazz during the birth of rock and roll, described how new musical styles really come about: "If you're going to come up with a new direction or a really new way to do something, you'll do it by just playing your stuff and letting it ride. The real innovators did their innovating by just being themselves." The funny thing is that not only was it completely unpredictable in the early twentieth century that jazz and blues would lead to rock and roll, but no one was even worried about it because *rock and roll was not an objective*. We just happened to be traveling the right path through the search space of musical genres to bump into it in the late 1940s [6].

One of the biggest names in popularizing rock and roll is of course Elvis Presley. Interestingly, his unmistakable sound wasn't planned either. Guitarist Scotty Moore recalls, "All of a sudden, Elvis just started singing this song, jumping around and acting the fool, and then Bill picked up his bass, and he started acting the fool, too, and I started playing with them. . . . [the recording engineer] stuck his head out and said, 'What are you doing?' And we said, 'We don't know [6].' " Who would have thought that "acting the fool," and not some ambitious drive to overhaul popular music, would transform the world of rock and roll? The stories of Elvis and rock and roll in general illustrate that not only can objectives block discovery, but that *having no objective* can lead to the greatest discoveries of all.

However, the problem is that it's hard to simply abandon objectives because they are a powerful security blanket. At the least they *seem* to protect us from the wild unknowns of the world. They give us a sense of purpose and the promise of success if only we try hard enough. No doubt, the thought of aimlessly wandering through the space of possibilities with no clear purpose is not going to inspire many can-do achievers. But that is not where this argument leads. We don't face a false choice between slavishly following objectives and aimless wandering. Instead, the true implications are both more subtle and more liberating than that. We want to show you that it's possible to explore a search space intelligently even without an objective. In other words, there *is* a third way—just because you don't have an objective doesn't mean you have to be wandering. We can align ourselves towards discovery and away from the trap of preconceived results.

Over the remainder of this book an important principle will begin to emerge: Sometimes the best way to achieve something great is to stop trying to achieve a *particular* great thing. In other words, greatness is possible if you are willing to stop demanding what that greatness should be. While it seems like discussing objectives leads to one paradox after another, this idea really should make sense. Aren't the greatest moments and epiphanies in life so often unexpected and unplanned? Serendipity can play an outsize role in life. There's a reason for this pattern. Even though serendipity is often portrayed as a happy *accident*, maybe it's not always so accidental after all. In fact, as we'll show, there's a lot we can do to attract serendipity, aside from simply betting on random luck.

We hope you'll find this message liberating. Our world has become saturated with objectives and metrics for success that mechanize our lives and distract us from our passions. But there are other paths to happiness and success. We'll show you why not only can you trust your gut instinct when it tells you something important is around the corner, but you *should* trust it, even if you can't explain what that something is. You don't need to make up a tortured *reason* to justify every little impulse you feel. And not only is this attitude more healthy for us as humans anyway, but it's backed up by solid scientific evidence that this book will present. We're missing out on a lot by clinging to objectives.

The objective culture around us is not natural anyway. How many children do you know who formulate an objective before they go out to play? How many great scientists really formulated a hypothesis *before* their great idea? Do you ever tell yourself that you can't do something because it's not *justified* by a clear purpose?

Do you ever blame yourself when you can't achieve your highest goals despite your best efforts? Why do so many of us feel our creativity is stifled by the machine-like integration of modern society? We weren't born for this. It isn't healthy for us, and deep down we know that. Something is wrong at a fundamental level with the assumptions that justify where our time is invested.

At the same time, this book is not only about personal liberation. In fact, objectives have become so pervasive that their absence would impact nearly everything. Their absence can change how science is conducted, how engineers conceive their projects, how architects generate new concepts, and how designers seek success. It can alter our understanding of natural evolution and how computer scientists develop algorithms. It can turn engineering into art and art into science. It can build bridges between disciplines and break down walls between others. It can redefine entrepreneurship and refocus our targets for investment. It can teach humility to the overconfident and confidence to the insecure. In short, objectives are a pillar of our culture, but they're also a prison around our potential. It's time to break out and discover what's outside.

Both of us, Ken and Joel, are computer scientists by training. Neither of us originally set out to attack or even question a topic so universal and seemingly solid as the benefits of objectives. Instead, we were (and still are) interested in artificial intelligence (AI), which was our primary area of expertise. However, it's funny how you never know where your path might lead, just like the theme of this book. It turns out that even in the field of AI objectives are everywhere. Many of the computer programs that AI researchers develop are designed to follow one objective or another, almost as if they're mimicking our own culture's drive to define success in the same way. For example, the objective for an AI program might be to help a robot learn to travel through a maze. With the field of AI itself so immersed in objectives, perhaps it was inevitable at some point that researchers in AI would eventually have to confront the downside of objectives themselves.

For us, the realization that there is a deep and largely unrecognized problem with objectives came from a particular AI project that we'll talk about in the third chapter. But the problem is a lot bigger than just about AI. It's about how you choose your major, find your spouse, and choose your profession. It's about how we make great discoveries, and how we pass our free time. It relates to happiness and success, but it's not even just about those things. It's also about understanding what really leads to major advances and major successes. Because once we understand that, our perspective can really change, and maybe then we can all spend more time on what's important and less time worrying about the myth—*the myth of the objective*.

Chapter 2
Victory for the Aimless

Life is what happens to you while you're busy making other plans.

John Lennon

Sometimes things work out the way we intended. Sometimes you study accounting and become an accountant. Or maybe you play basketball as a kid and make it onto the varsity team. But the great successes, the ones that come crashing out of nowhere and shake up the system, they don't usually follow this kind of script. You don't enroll in Superstar 101 to become a superstar. There's no magic formula for changing the world. In other words, the greatest victories are not written into the initial plans. They happen despite the plans.

Take career choice for example. There's no shortage of resources to help you choose your target career. You can read Richard Bolles' *What Color Is Your Parachute?* to find out how to choose and secure the job that's right for you [7]. Or you could take the *Career Key* test to discover what you really should be doing with your life [8]. The list is endless. There's the *Campbell Interest and Skill Survey* [9], personality tests like the *Myers-Briggs Type Indicator* [10] and *Keirsey Temperament Sorter* [11], skillset assessments, work values tests, and so on. Amazon.com shows over 97,000 books on careers. If you can't find your way, there's always a guru eager to guide you.

But while a little guidance may be good for some, others don't follow the script. Some people set objectives and completely miss them—all for the better. Why is this kind of story so common for the most successful? Why don't they need to stop in at a career counselor? At heart, it's not such a mystery. It's just that anticipating what might lead to the most fulfilling outcomes is difficult. As with all open-ended problems in life, the stepping stones are unknown. So when you go out into the uncertain world sometimes it may be wise to hitch a ride with serendipity. Being open and flexible to opportunity is sometimes more important than knowing what you're trying to do. After all, any path might lead to happiness, even the most unexpected. Some people seem to have an uncanny knack for spotting such opportunities, even if they conflict outright with their original aims.

K.O. Stanley, J. Lehman, *Why Greatness Cannot Be Planned*,
DOI 10.1007/978-3-319-15524-1_2

In an interesting experiment by Richard Wiseman [12], subjects were asked to count the number of photographs in a newspaper. It turns out that those who focused on the goal of counting the photographs took significantly *longer* to complete the task than those who were less focused on the objective. Why? The more open-minded participants noticed that on the inside of page two Wiseman had written, "Stop counting: There are 43 photographs in this newspaper." While some might say that noticing the answer on page two is only luck, the deeper lesson is that focusing too much on your goal can actually prevent you from making useful unexpected discoveries.

Some might also say Johnny Depp was just lucky. Who would have guessed that being in a mildly successful band would somehow lead to an incredibly successful career in acting? Wouldn't it be smarter to fill your schedule with hours of acting classes? But starting a band was the right stepping stone for Johnny Depp, though not so much because it was logical or planned. It was simply that Depp wanted to make music and was open to opportunities when they came. In fact, even his high school principal once told him to pursue his dream of musical stardom instead of finishing high school. While his music career never blossomed, it turns out that joining a band had some unexpected benefits. Not only did Depp marry the bass player's sister, but she ultimately introduced him to acting through her job as a makeup artist [13]. If Johnny Depp had not wanted to become a musician, he may never have become an actor.

For the very successful, these kinds of stories are surprisingly common. Before he wrote books, the bestselling novelist John Grisham first trained and practiced as a criminal defense attorney for ten years. The trigger for his career change was particular testimony that he overheard one day from a young rape victim. Somehow that testimony made him realize that he should and could write, and he began waking early in the morning before work to gradually complete his first novel, *A Time to Kill* [14]. His next book, *The Firm*, would spend 47 weeks on the New York Times bestseller list. Most aspiring writers would not choose law school to develop their craft. Reading endless case studies in a dusty library seems like poor preparation compared to practicing creative writing. But maybe that's why it worked for John Grisham. He wasn't following a plan.

In fact, that seems to be a pretty good strategy if you want to be a writer. Before *Harry Potter* sold millions of copies, J.K. Rowling was a bilingual secretary for Amnesty International and then taught English to students learning it as a second language in Portugal [15]. Haruki Murakami, the Japanese writer behind such award-winning reveries as *The Wind-Up Bird Chronicle* and *Kafka on the Shore*, first ran a combination coffee house and jazz bar. Murakami himself is convinced that without running the bar he never would have become a writer because it gave him the time to observe and brood. Many of his characters later shared his fondness for jazz. Interestingly, it didn't occur to Murakami that he could write novels until he was 29 [16].

And the examples keep coming. The founder of American hard-boiled detective fiction, Raymond Chandler, didn't write until he was fired from his job as an oil company executive at age 45 [17]. And as late-blooming philosopher Mary Midgley

that future President Coolidge ever thought his best moment would involve being caught almost naked. But then again, the unplanned is often the best plan of all. At least in the domain of romance, we all have some experience with the paradox of ambitious objectives. After all, what could be more ambitious than seeking lifelong happiness?

Just as that lesson is familiar for love, it's also sometimes present in other spheres, like our recreations. For example, unlike with careers we often choose hobbies not because of some long-term grandiose plan, but just because we like them. So again having no objective leads at least to some personal satisfaction. The Internet has made it easier now to share and discover odd hobbies like "snail racing," "underwater hockey," "limbo skating," "extreme unicycling," and even "extreme ironing" (all of which you can find on Wikipedia). Some hobbies do even become stepping stones to something bigger—Nathan Sawaya was a corporate attorney who had so much fun building artistic models out of Lego blocks that he quit his job to do it full time [27]. But despite the risks of deviating from his original plan, following his heart in this case led again to something better and now his unique pieces sell well enough to make a good living.

Or maybe you've heard of Joseph Herscher, who spends much of his time building intricate Rube Goldberg machines with no real necessity other than being fun to watch [28]. In one contraption, a ball rolls down a ramp and falls onto a lever that lights a fuse, burning away a string that releases another ball and so on, all ultimately accomplishing nothing but opening his newspaper. It might sound pointless (a popular word for having no clear objective) but as usual there are considerable benefits: His work has been viewed by millions, leading to numerous television and radio appearances. Indeed, psychologists have noted that children need time to explore without specific tasks or objectives set for them by adults [29, 30]. Sometimes the term *unstructured play* is used to describe this kind of activity. Maybe adults need it too.

Of course, when you hear about people spending their adult lives playing with Legos or building Rube Goldberg machines, it's easy to dismiss their passions as frivolous. But there's a deeper value to such endeavors than just the endeavors themselves: They reflect that we don't know which stepping stones might lead to something interesting. These are people who are willing to commit their lives to stepping stones that most of us would entirely ignore, which is good for all of us. Because no one knows the stepping stones that lead to the greatest discoveries, the last thing we want to do is stop people from exploring stepping stones that we ourselves choose to ignore—who knows what they will find? It's not that we are all one day going to wake up desperately seeking Rube Goldberg machines. But what if Joseph Herscher one day wakes and realizes that one of his Rube Goldberg machines solves an unexpected problem?

A historical example of just this kind is given by the unexpected discovery of ancient cave paintings by a non-scientist, Marcelino de Sautuola, near the hill of Altamira in Spain in 1879 [31]. Before this discovery, people had no idea of how sophisticated prehistoric paintings could be. Importantly, Marcelino held a wide variety of unique hobbies, including caving and learning about ancient artifacts.


image
not
available