Online Version

Harvard Business Review

THE HBR INTERVIEW

The Institutional Yes

How Amazon's CEO leads strategic change in a culture obsessed with today's customer.

An Interview with Jeff Bezos by Julia Kirby and Thomas A. Stewart

Amazon.com was born of strategy. The story has often been told of how founder Jeff Bezos, working in a quantitative analysis group at an investment firm, spotted an opportunity to sell books on the Internet. No native of the book-selling industry, he arrived at his business model logically: Given the attributes of the product and the structure of the supply chain, a no-bricks retailer could clearly make it—and make it big. Since its founding, in 1995, Amazon has continued to show a knack for spotting white spaces and a willingness to jump into them, even as it works to make spaces it already occupies more productive. The company's latest bold move is its "developer-facing" business: a set of offerings that make tools Amazon developed for its own use available to other website developers. Like so many other Amazon ventures, this one is based on sound logic—yet it's surprising. It's not the kind of thing one sees many companies doing.

Enough of these moves have paid off that HBR decided to look into what's different about strategy formulation at Amazon. How do ideas come under consideration, and how are commitments made? Is it a matter of one ingredient—Jeff Bezos—or is it an institutional capability? Thomas A. Stewart, HBR's editor in chief, and Julia Kirby, a senior editor, met with Bezos on two occasions to learn more. The conversations were wide-ranging and lively, punctuated frequently by Bezos's famous laugh. (He says that his wife, MacKenzie, often tells people, "If Jeff is unhappy, wait five minutes.") They came away with a sense that Amazon's strategy and culture are rooted in a sturdy entrepreneurial optimism. The following is an edited presentation of those talks.

Who is setting strategic direction for Amazon? At the very beginning it was just you, sitting in a car on the way from New York to Seattle, making all the plans. Are you still making them all?

Oh, heavens, no. We have a group called the S Team—S meaning "senior"—that stays abreast of what the company is working on and delves into strategy issues. It meets for about four hours every Tuesday. Once or twice a year the S Team also gets together in a two-day meeting where different ideas are explored. Homework is assigned ahead of time. A lot of the things discussed in those meetings are not that urgent—we're a few years out and can really think and talk about them at length. Eventually we have to choose just a couple of things, if they're big, and make bets.

The key is to ensure that this happens fractally, too, not just at the top. The guy who leads Fulfillment by Amazon, which is the web service we provide to let people use our fulfillment center network as a big computer peripheral, is making sure the strategic thinking happens for that business in a similar way. At different scale levels it's happening everywhere in the company. And the most important thing is that all of it is informed by a cultural point of view. There's a great Alan Kay quote: "Perspective is worth 80 IQ points." Some of our strategic capability comes from that.

How would you describe that cultural point of view?

First, we are willing to plant seeds and wait a long time for them to turn into trees. I'm very proud of this piece of our culture, because I think it is somewhat rare. We're not always asking ourselves what's going to happen in the next quarter, and focusing on optics, and doing those other things that make it very difficult for some publicly traded companies to have the right strategy.

Do you know when you're planting one of those seeds that it's, say, an acorn and it's going to turn into an oak? Do you have a strong vision of how things will materialize? Or does the shape emerge along the way?

We may not know that it's going to turn into an *oak*, but at least we know that it can turn out to be that big. I think you need to make sure with the things you choose that you are able to say, "If we can get this to work, it will be big." An important question to ask is, "Is it big enough to be meaningful to the company as a whole if we're very successful?"

Every new business we've ever engaged in has initially been seen as a distraction by people externally, and sometimes even internally. They'll say, "Why are you expanding outside of media products? Why are you going international? Why are you entering the marketplace business with third-party sellers?" We're getting it now with our new infrastructure web services: "Why take on this new set of developer customers?" These are fair questions. There's nothing wrong with asking them. But they all have at their heart one of the reasons that it's so difficult for incumbent companies to pursue new initiatives. It's because even if they are wild successes, they have no meaningful impact on the company's economics for years. What I have found—and this is an empirical observation; I see no reason why it *should* be the case, but it tends to be—is that when we plant a seed, it tends to take five to seven years before it has a meaningful impact on the economics of the company.

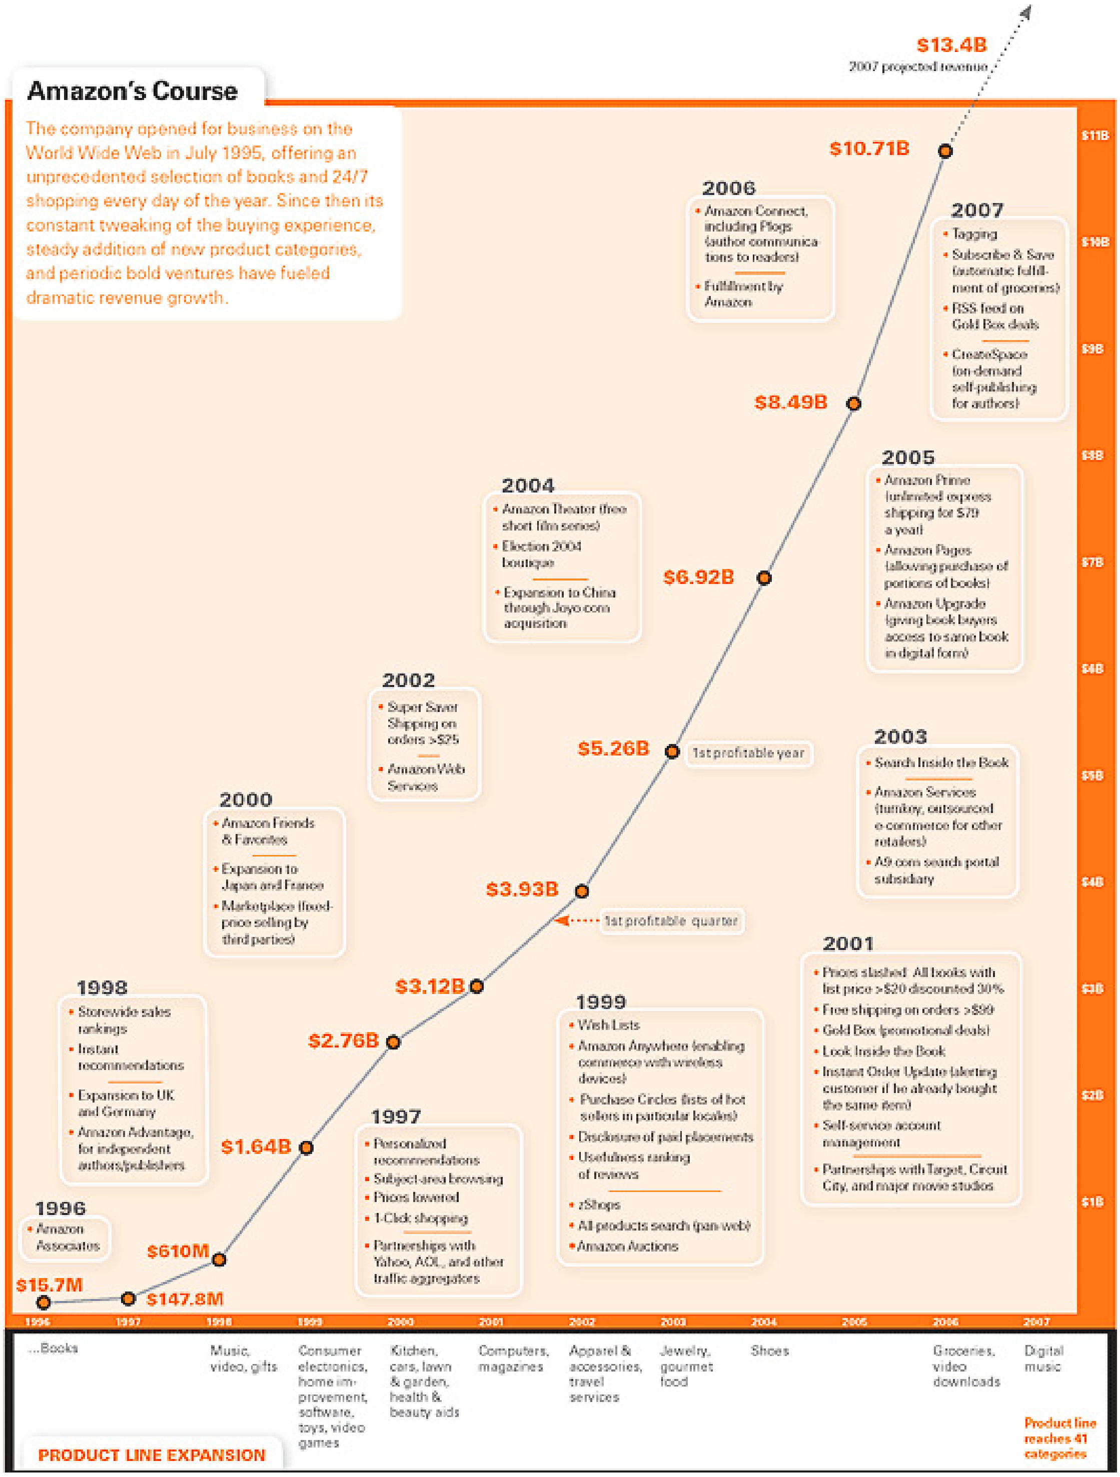

That does require people, inside and outside, to keep the faith. How do you have the confidence that the investment will ultimately pay off?

It helps to base your strategy on things that won't change. When I'm talking with people outside the company, there's a question that comes up very commonly: "What's going to change in the next five to ten years?" But I very rarely get asked "What's *not* going to change in the next five to ten years?" At Amazon we're always trying to figure that out, because you can really spin up flywheels around those things. All the energy you invest in them today will still be paying you dividends ten years from now. Whereas if you base your strategy first and foremost on more transitory things—who your competitors are, what kind of technologies are available, and so on—those things are going to change so rapidly that you're going to have to change your strategy very rapidly, too.

What are some of the things you're counting on not to change?

For our business, most of them turn out to be customer insights. Look at what's important to the customers in our consumer-facing business. They want selection, low prices, and fast delivery. This can be different from business to business: There are companies serving other customers who wouldn't put price, for example, in that set. But having found out what those things are for our customers, I can't imagine that ten years from now they are going to say, "I love Amazon, but if only they could deliver my products a little more slowly." And they're not going to, ten years from now, say, "I really love Amazon, but I wish their prices were a little higher." So we know that when we put energy into defect reduction, which reduces our cost structure and thereby allows lower prices, that will be paying us dividends ten years from now. If we keep putting energy into that flywheel, ten years from now it'll be spinning faster and faster.

Another thing that we believe is pretty fundamental is that the world is getting increasingly transparent—that information perfection is on the rise. If you believe that, it becomes strategically smart to align yourself with the customer. You think about marketing differently. If in the old world you devoted 30% of your attention to building a great service and 70% of your attention to shouting about it, in the new world that inverts. A lot of our strategy comes from having very deep points of view about things like this, believing that they are going to be stable over time, and making sure our activities line up with them. Of course there could also come a day when one of those things turns out to be wrong. So it's important to have some kind of mechanism to figure out if you're wrong about a deeply held precept.

You've got two other sets of customers: the third parties who are selling goods through your site and this new set—the developers who can benefit from the tools you've created over the years. What are some constants for those groups?

We're still working on identifying them for the developer community, although we have some good guesses as to what they are. Reliability of the platform would be one, which is kind of a no-brainer. But then a lot of these things are no-brainers. No-brainers are no-brainers for a reason: They actually are important. As for the sellers, the number one thing that sellers want is sales.

Is that why the auction business didn't work out for Amazon—because eBay already had a lock on the sales those sellers wanted?

Actually, no, it's a little different from that. It's that our customers who are buyers are very convenience motivated. We make it really, really easy to buy things. You can see that if you look at a metric like our revenue per click or revenue per page turn. It's very high, because we're efficient for people. If you're a customer who wants that kind of quick service, you do not want to wait till an auction closes. An auction is more about playing a game. There's some fun involved. You're not necessarily just trying to get the job done. It's a different kind of thing and a different customer segment.

That episode is actually one of the highlights of our corporate history—one that I tell over and over internally, because it speaks to persistence and relentlessness. The basic thought was: Look, we have this website where we sell things, and we want to have vast selection. One of the ways to get vast selection is to invite other sellers, third parties, onto our website to participate alongside us, and make it into a win-win situation. So we did auctions, but we didn't like the results. Next we created zShops, which was fixed-price selling but still parked those third parties in separate parts of the store. If a third-party seller had a used copy of *Harry Potter* to sell, it would have its own detail page, rather than having its availability listed right next to the new book. We still didn't like the results we got. It was when we went to the single-detail-page model that our third-party business really took off. Now, if we're offering a certain digital camera and you're a seller with the same camera to sell, you can go right on our own detail page, right next to our product, and underbid us. And if you do, we will put you in the "buy" box, which is on that page.

That can't have been an easy decision—it gets you the seller customer but loses you the buyer customer, doesn't it?

It was a very controversial decision internally at the time. Imagine being our digital camera buyer and you've just bought 10,000 units of a particular digital camera. Your boss says to you, "Good news: You know all those people you've been thinking of as your competitors? We're going to invite them to put their digital cameras right next to yours on your detail page." The natural reaction is *"What?"* A digital camera today is like a fish in a supermarket. It's aging fast. You don't want to get stuck with a huge inventory of five-megapixel cameras when customers will shortly want six-megapixel cameras. You'll have to sell them for pennies on the dollar. So our buyers were extremely concerned—and rightly. They were saying, "Let me just make sure I understand this. I might get stuck with inventory of 10,000 units of this camera that I just loaded up on, and you're going to let just anybody come in and take Amazon traffic on what is our primary retail real estate, which is the detail page, and I'm going to lose the buy box to this other person because they have a lower price than me?" And we said, "Yeah, we are."

We talked about that a lot before we did it. But when the intellectual conversation gets too hard because of these potential cannibalization issues, we take a simpleminded approach. There's an old Warren Buffett story, that he has three boxes on his desk: in-box, out-box, and too hard. Whenever we're facing one of those too-hard problems, where we get into an infinite loop and can't decide what to do, we try to convert it into a straightforward problem by saying, "Well, what's better for the consumer?"

The reality is that we have a highly competitive cost structure and we're very good at competing on things like digital cameras on that page, so it actually works out very well. But some of the most important things we've done over the years have been short-term tactical losers. In the very earliest days (I'm taking you back to 1995), when we started posting customer reviews, a customer might trash a book and the publisher wouldn't like it. I would get letters from publishers saying, "Why do you allow negative reviews on your website? Why don't you just show the positive reviews?" One letter in particular said, "Maybe you don't understand your business. You make money when you *sell* things." But I thought to myself, We don't make money when we sell things; we make money when we help customers make purchase decisions.

So you're choosing projects based on stable customer needs and adopting the customer's perspective to make difficult strategic decisions. Michael Porter claims that strategy is shaped by five forces, but we're hearing you talk about only one.

I think somehow I am congenitally customer focused. And I think that from that comes this passion to figure out customer-focused strategies as opposed to, say, competitor-focused strategies. There's nothing wrong with competitor focus. There are other companies that have delivered exceedingly good business results by pursuing close-follower strategies, setting up really good benchmarking tools. There's a lot to be said for that kind of passion. It doesn't happen to be who we are. We don't ignore our competitors; we try to stay alert to what they are doing, and certainly there are things that we benchmark very carefully. But a lot of our energy and drive as a company, as a culture, comes from trying to build these customer-focused strategies. And actually I do think they work better in fast-changing environments, for two reasons. First, as I said before, customer needs change more slowly—assuming you pick the right ones—than a lot of other things. Second, close following doesn't work as well in a fast-changing environment. The strategic value of close following is in not having to go down all the blind alleys. You let smaller competitors check those out, and when they find something good, you just quadruple down. If you're following close enough, and the arena is slow-moving enough, the fact that you're not first down that path doesn't hurt you much. But in our environment there's so much rapid change on the Internet, in technology, that our customer-obsessed approach is very effective.

There's another situation, too, where I think being customer focused is better, and that's when you're a leader. If you're competitor focused, you tend to slack off when your benchmarks say that you're the best. But if your focus is on customers, you keep improving. So there are a lot of advantages. I'm not claiming we invented this approach—a lot of companies are customer focused—but it's very deeply ingrained in all the nooks and crannies of our culture.

It could be that you have been able to take it to a new level. Being the instantiation of that model at themoment when it became possible to get a flood of customer feedback was happy timing.

It's totally true. We have the opportunity for Amazon not just to be a customer-centric company but to set a new standard globally for what "customer-centric" means. After World War II, Morita-san [Akio Morita, Sony's longtime leader] set a goal for Sony. He wanted the company to be known for quality, but his bigger goal for Sony was to make Japan known for quality. Having that kind of bigger mission is very inspiring. Years from now, when people look back at Amazon, I want them to say that we uplifted customer-centricity across the entire business world. If we can do that, it will be really cool.

In many companies it's hard to keep that customer focus. As the enterprise itself grows, internal agendas start to drive decisions. How do you defend against that? Are you constantly broadcasting the points of view that you hope will deeply inform the culture?

We have three all-hands meetings a year, so that's one mechanism for it. And then, repetition-wise, we've been saying the same things for many years. I constantly talk about things like information perfection and customer obsession relative to competitor obsession. We also do a bunch of things to keep people directly in touch with customer needs. Every new employee, no matter how senior or junior, has to go spend time in our fulfillment centers within the first year of employment. Every two years they do two days of customer service. Everyone has to be able to work in a call center.

Even you?

Oh, yeah. I just got recertified about six months ago. The fact that I did a lot of customer service in the first two years has not exempted me. Besides, it's quite entertaining, and you learn a ton. It's not a chore.

These mechanical things continually add fuel to that preexisting fire, but the truth is that corporate cultures are incredibly stable over time. They are self-perpetuating, because they attract new people who like that kind of culture, while the people who don't like it eject themselves. One of the things I hear most from new executives, once they've been here for a month, is "I can't believe how customer focused this place is." So if you were someone who really loved a competitor-obsessed company and you came to Amazon, one of two things would happen. Either you would see our approach and say, "Wow, this is actually pretty cool; my eyes have been opened," or you would say, "This isn't for me." You might not use the word "culture," but you'd be missing something. So there's a very strong self-reinforcing loop. That is why cultures are very hard for competitors to copy—and a source of great competitive advantage if you have the right culture for the right mission. They also limit. There are things that you cannot undertake. You have to be accepting of what your culture is in your thinking about strategy.

That would mean the culture of a company is very path dependent. With the first few people you hire and your first few experiences, things start down a path.

They absolutely do. I'll tell you a story about an incident that I'm sure is part of the reason we are what we are. Sometime around May of 1997 we were put on a deathwatch by a well-known industry pundit. His argument was a simple one: Barnes & Noble had just launched their website, and we had only 125 employees at the time, and something like $60 million in annual sales.

So his point was that you'd awakened a sleeping giant and filled it with a fierce resolve?

Exactly. They had 30,000 employees at that time, and $3 billion in sales. His message was, Amazon has had a good two-year run, but it's over. They're toast. The title of his report was actually "Amazon.Toast." It got picked up by the media, and it was pretty much everywhere, because we were sort of a poster child Internet company. To a person, I think, every one of our people got called by a parent—usually their mother—asking, "Are you okay? What's this about Amazon dot Toast?"

We had an all-hands meeting—which was very easy to do with only 125 people—and said, "Forget about this. We can't be thinking about how Barnes & Noble has so much more in the way of resources than we do." I told everyone, "Yes, you should wake up every morning terrified with your sheets drenched in sweat, but not because you're afraid of our competitors. Be afraid of our customers, because those are the folks who have the money. Our competitors are never going to send us money."

Since then we've made a number of decisions siding with the customer that have been questioned by well-meaning critics, journalists, Wall Street analysts, and industry analysts. I'm talking about things like free shipping, relentlessly lowering prices, Amazon Prime. And every time we make one of those and get broadly criticized for it and then it works, we get a little more credibility.

What would you say has been the nature of your biggest strategic mistakes?

I think most big errors are errors of omission rather than errors of commission. They are the ones that companies never get held to account for—the times when they were in a position to notice something and act on it, had the skills and competencies or could have acquired them, and yet failed to do so. It's the opposite of sticking to your knitting: It's when you shouldn't have stuck to your knitting but you did.

It's hard to imagine how you could spot those errors of omission in time to hold a management team accountable. What are you doing to prevent them?

One useful habit is to ask the question, "Why not?" When something seems like an opportunity—it seems like you have the skills, and maybe some kind of advantage, and you think it's a big area—you will always get asked the question, "Why? Why do that?" But "Why not?" is an equally valid question. And there may be good reasons why not—maybe you don't have the capital resources, or parts of your current business require so much focus at this key juncture that it would be irresponsible. In that case, if somebody asked, "Why not?" you would say, "Here's why not…" But that question doesn't get asked. It's an asymmetry that is linked to those errors of omission.

If you are determined to avoid errors of omission, doesn't that put you at great risk of taking on too much and spreading resources too thin?

What you really want to do companywide is maximize the number of experiments you can do per given unit of time. If something's really big—like the big bet we've made on Amazon Web Services—then sure, you can do only a limited number of those, so you spend more time thinking about them and talking them through. Somebody wears the black hat and makes the case for why not to do it, and somebody else puts on the white hat and says why it is actually a good thing to do. But since the outcomes of all these things are uncertain, if you can figure out how to conduct an experiment, you can make more bets. So the key, really, is reducing the cost of the experiments. We have a group called Web Lab that is constantly experimenting with the user interface on the website, getting statistical data from real usage patterns about which interfaces work best. That is a huge laboratory for us, and we've put a lot of energy into trying to figure out how to be very low cost with those experiments so that we can run a much larger number of them.

As a practical matter, sometimes it's very, very hard to dramatically reduce the cost of experimentation. There are areas where conducting an experiment is still, in terms of cost, tantamount to just doing it. But you should always be trying to do that. You should be making guesses and then finding out at lowest cost whether or not they are needle movers.

How often do experiments turn out wildly different from what you expected?

Sometimes you make guesses and you think, When we launch this, people are going to love it. And they don't. We launched a feature—I forget what we named it—where we'd take your purchase history and then find, out of our millions of other customers, the one customer whose purchase history was the closest match. You could push a button, and we'd show you all the things that your doppelgänger had bought that you had not. No one used it. Our history is full of things like that, where we came up with an innovation that we thought was really cool, and the customers didn't care. Fortunately, there are also quite a few that went the other way. Back when we first launched the Associates program, we thought it was a nice thing to do but had relatively low expectations for it. This is our marketing program that lets other sites earn affiliate fees by sending buyers to us. It was the first program of its kind, and it dramatically exceeded our expectations. Very quickly we doubled down on it as a favored marketing program. And it's continuing to be very successful 11 years later.

At the same time, you have to recognize that there are times when you can't put a toe in the water; you have to leap in with both feet. You have to say, "This is going to be expensive—and that means we're going to have to make it work." You allow that you might take lots of twists and turns on the details, but you really commit yourself to the objective. And by the way, it's very fun to have the kind of culture where people are willing to take these leaps—it's the opposite of the "institutional no." It's the institutional yes. People say, "We're going to do this. We're going to figure out a way."

How does that square with the emphasis you place on taking your cues from the early results of experiments?

My observation on that would be that it's important to be stubborn on the vision and flexible on the details. I talked about the evolution of our marketplace business—that's a good example of where we were relentless on the vision. We made a lot of twists and turns in the execution. We worked on it for a few years. But we didn't give up on the vision. I think if you believe that a particular customer set is important, like sellers in that case, and you believe you have identified what some of their nontransient needs will be, and you believe the addressable market is big enough for it to matter—if you really believe those things, then it pays to be stubborn in pursuing that.

As the business grew and moved into new areas, you must have felt challenged to keep up on a personal level. What have you had to learn along the way?

Something we haven't talked about, but that is super important in our culture, is the focus on defect reduction and execution. It's one of the reasons that we have been successful for customers. That is something I had to learn about.

That's not your natural strength?

Well, by "learn" I mean I literally learned a bunch of techniques, like Six Sigma and lean manufacturing and other incredibly useful approaches. I'm very detail oriented by nature, so I have the right instincts to be an acceptable operator, but I didn't have the tools to create repeatable processes and to know where those processes made sense. Before I started Amazon, I was with a quantitative hedge fund. That work is very disciplined and very analytical, but it's not a question of designing a repeatable process. It's not like a car-manufacturing plant, where the work has to be done in a defect-free way, in the same way over and over and over—or anyway, that piece of it is not as important. But here, that execution focus is a big factor, and you can see it in our financial metrics over the past ten years. It's very obvious when, for instance, we look at the number of customer contacts per unit sold. Our customers don't contact us unless something's wrong, so we want that number to move down—and it has gone down every year for 12 years. That's big-time process management. We try to implement those kinds of processes in various places. They're most naturally applied in our fulfillment centers and in customer service and so on, but we've actually found that they can be useful in a bunch of different things.

When you are inexperienced with disciplined process management, you initially think that it's equivalent to bureaucracy. But effective process is not bureaucracy. Bureaucracy is senseless processing—and we've had some of that, too.

What else has been critical to your ability to grow as a manager and leader? Often, entrepreneurs aren't able to scale with the businesses they found.

I don't know that I have anything profound to say about that. When you start out, it's a one-person thing, at least on the first day, and you're not only figuring out what to do but actually doing it. At a certain point the company gets bigger, and you get to where you're mostly figuring out what to do but not how it's done. Eventually you get to the point where you're mostly figuring out who is going to do it, not even what to do. So one way to think about this is as a transition of questions, from "How?" to "What?" to "Who?" It's not necessarily a conscious thing; circumstances dictate it. As things get bigger, I don't think you can operate any other way. So the question for any senior executive is, basically, "Are you progressing along that curve at the right rate?"

Now, of course, I'm oversimplifying, because the reality is that you continue to dive deep in selected areas. I've never met what I think of as a good executive who does not choose certain highly leveraged activities—some area they consider so important that they will inspect it all the way down to the how stage.

What's your deep-dive area?

When it comes to the way we relentlessly drive down our consumer-facing pricing, I still continually launder and inspect that and talk to the people who do the work all the way through that whole chain. I need to be sure that we are in fact competitive and focused on offering our customers the lowest possible prices. That's one of the things I think is so highly leveraged that I am involved from heading level one all the way to heading level five.

Is there anyone in the company who can say no to you?

I'm actually thinking, Who doesn't? We have many strong executives, so not even just my direct reports but beneath them. We have an informal atmosphere, which I think helps people tell me no—and not just me. It's also really important that they be able to say what they think to their senior vice president or vice president and so on. An informal atmosphere, I think, is a huge benefit.

That doesn't mean we don't have raucous debates and get angry with each other occasionally. Intensity is important. I always tell people that our culture is friendly and intense, but if push comes to shove, we'll settle for intense. But there is no contradiction between being intense and having fun. You can absolutely do that. That's what we try to achieve, and at our best it is what we achieve. We'll have some big problems, we'll get together, and we'll laugh about them. Sometimes it's dark humor, but it's either laugh or cry. And then we try to work our way through them. I think there's a lot of value in having fun, even though in a more formal culture some of our meetings might look undisciplined. In a one-hour meeting we may spend ten minutes of it joking around, and I'm often the worst offender. I'll laugh and say, "This reminds me of…" and get us off on some story. Eventually somebody says, "Well, that is very interesting, but you do see we have an agenda…" And I think that works out great.