



Dear Shareholders:

When I graduated from college, I wanted to be a sportscaster. After sending my resume reel to many small markets around the U.S., and only getting two nibbles, I settled on doing sports production at a major network. To make extra money, I also coached my former high school soccer team, and worked at a retail golf store. Six months later, a college classmate convinced me to interview at the consumer products company where he worked, and I spent three years as a Product Manager there. I left that job to try some of my own businesses. After deciding these businesses weren't my calling, I tried short stints in sales and investment banking, before going back to graduate school and ending up at Amazon three days after my last final exam in May 1997.

Not exactly a straight line.

AWS followed lots of squiggly lines, too. The original vision included storage, compute, payments, and human intelligence. Some of those (e.g. storage and compute) became lynchpins in AWS. Others didn't succeed. We didn't initially plan a database service; and when we built one, our first attempt failed to get traction. We went back to the drawing board and built new relational and non-relational database services, which have resonated well and become core to millions of AWS applications. When we launched EC2 (our compute service), it was a single instance type in one availability zone, Linux-only, with no auto-scaling, load balancing, block storage, or private networking. Over time, we added those capabilities and hundreds more services. AWS was initially attractive to start-ups (companies like DoorDash, Dropbox, Pinterest, Slack, and Stripe were among many that built their businesses on AWS). Pundits said enterprises and governments would never use cloud or AWS for anything substantive. In 2008, Netflix decided to move all of its applications to AWS, then came big commitments from GE, Intuit, and others—and eventually the CIA chose AWS as their partner to build their classified cloud. Growth came fast and furious, and as it accelerated, so too did our capital expenditures (“capex”) with dilutive impact on free cash flow (“FCF”). At our 2014 AWS operating plan review, the discussion started with a senior leader at the company musing, “Tell me again why we're doing this business?”

AWS has worked out well for Amazon, but a straight line? Not really.

There's a band I like from New Zealand called “The Beths,” who've written several excellent records, with thought-provoking lyrics. I eagerly await their new releases, and when their latest album dropped last summer titled “Straight Line Was a Lie,” it made me think about how prescient that expression is. Most long-term endeavors do not follow a linear straight line, up and to the right. Progress jumps around; it'll zig up, then sometimes stall, or zag down, or force you back to the starting line. Sometimes, it feels like you're running in circles. But, the path is rarely straight.

That's because the world is complex, and new technology, business model invention, competitors, global issues, or people and cultural shifts can come into play. We're in the middle of some of the biggest inflections of our lifetime (e.g. AI, robotics, space industrialization, geopolitical and military conflict). And, just as proficient golfers need to be skilled across driving, approach shots, chipping, and putting, durable companies must be adept at managing different elements of inflections. I'll share some of our lessons below and why we're bullish on what's ahead for Amazon.

Wherever possible, invent the next inflections. We try to anticipate what will make customers' lives easier and better every day, and invent the next inflection. Historically, we've successfully done so in areas like Retail, Logistics, AWS, Ads, Kindle, Alexa, and Pharmacy. There are too many new efforts in flight to mention them all, but will mention a few.

First, despite many improvements over the years, customers always want lower costs and faster delivery speed. While we continue to work on productivity and inventory levels, robotics provides a step-level change

for how we can deliver faster, reduce the cost of carrying more selection, and automate movements that cause strains and injuries to our teammates. Accelerated by acquiring Kiva in 2012, and investing in numerous robotics initiatives the last 14 years, we now have over one million robots operating in fulfillment centers helping with stowing, picking, sorting, and intra-facility transport. And, we've done this while continuing to be one of the largest job creators in the country. While the above progress is substantial, we're still in the early stages of how we'll leverage robotics. Expect us to keep innovating on form factors, use case diversity, agility, grasping, and intelligence. And, wherever we can leverage our scale and real-time feedback loop from so many robots in our fulfillment network to build robotics solutions for other industrial and consumer customers, we'll explore doing so.

Second, we understand that rural customers are often de-prioritized by logistics and telecom providers because remote communities are more expensive to serve. While other companies have been backing away from these customers, we've been running to them. We've committed over \$4 billion to expand our rural delivery network. Customer response has been overwhelmingly positive, with the average number of monthly Same-Day customers in rural areas nearly doubling in 2025 compared to the prior year. Once this expansion is complete, our network will be able to deliver over a billion more packages each year to customers living in over 13,000 zip codes spanning 1,200,000 square miles.

We're also trying to close the digital divide for rural communities. There are billions of people on the planet who lack high-speed internet access, and millions of businesses, governments, and other organizations operating in places without reliable connectivity. If you don't have broadband connectivity, you can't engage in many of the digital activities (e.g. education, business, information retrieval, shopping, entertainment, etc.) that people take for granted in metropolitan locations. Over the last seven years, we've built a low Earth orbit satellite network (Amazon Leo) and put more than 200 satellites into space (which is the third-largest low Earth orbit network operating today). With a few thousand more satellites launching in the coming years, the constellation is expanding rapidly. Apart from enabling this connectivity, Leo will offer three unique benefits. First, the performance will be stronger (about six to eight times better on uplink, and two times better on downlink) than what customers have access to now. Second, this performance will come at a lower cost than alternatives. And third, Leo will seamlessly integrate with AWS to enable enterprises and governments to move data back and forth for storage, analytics, and AI.

While Amazon Leo is officially scheduled to launch in mid-2026, we already have meaningful revenue commitments from enterprises and governments. Most recently, Delta Airlines, the highest grossing airline in the world, has announced it's chosen Amazon Leo for its future Wi-Fi, and will begin with 500 planes in 2028. They join other Leo customers like JetBlue, AT&T, Vodafone, DIRECTV Latin America, Australia's National Broadband Network, NASA, and others.

Amazon could be successful for a long time without investing this way in robotics, faster rural delivery, and broadband connectivity for underserved customers and geographies. But, we believe we can invent ways to change what's possible for customers, are hungry to do so, and are confident these investments will yield meaningful growth and return on invested capital ("ROIC") for the company.

Be willing to pursue parallel paths when it's unclear what'll best drive the desired trajectory (2 > 0). When I was a kid, I used to go to New York Rangers games with my Dad. I loved hockey, and it was high quality time together. I looked up to my Dad (still do), and hung on his every word. One game, my Dad noticed that one of the Rangers' defensemen, Dallas Smith, had gone missing from the bench, and stood up and exclaimed "where's Dallas?" to which a nearby fan said, "in Texas, moron."

So yes, it's fairly obvious 2 is greater than zero. But too often, companies focus on what looks most tidy instead of ensuring they have enough efforts in play to achieve an important outcome. Let's go back to fast delivery in our retail business. We know how much customers crave it, and we see higher order completion rates when delivery promises are faster. Just three years ago, two-day delivery was the gold standard. We pushed that bar to one day, and have been working tirelessly to make it same day.

We've invented a new, more streamlined fulfillment center format called Same Day Fulfillment Centers ("SSDs"). We've built over 85 SSDs across the U.S. that carry our top 90,000 SKUs, and have enabled us to deliver more than 500 million same day units in 2026 thus far. At the same time, we've continued to

pursue Prime Air, our drone delivery service. Prime Air now has a design that'll scale, plans to serve communities with 30 million customers by year-end, and expects to deliver half a billion packages by the end of this decade (with an aim to deliver inside 30 minutes). And, over the last year, starting in India and the UAE, we've been working on Amazon Now, ultra-fast delivery on thousands of items within 20 minutes. Customers love this service. In India, where we have more than 360 micro-fulfillment centers (and growing rapidly), Amazon Now orders are increasing 25% month-over-month, with Prime members tripling their shopping frequency once they start using it. We're starting to expand Amazon Now in the U.S. and Europe, too.

Some companies may have decided to pursue only one of these efforts, arguing for weeks or months which one, all the while pursuing none. A reasonable argument could have been made that just delivering same day would have been enough ("get that done and worry about the other ideas in the future"). But, building an autonomous delivery drone that can deliver millions of items in 30 minutes doesn't take a year. It's a multi-year invention cycle. And, ultra-fast delivery from micro-fulfillment centers is already possible and going to happen with or without us. So, we need to have multiple parallel paths to drive this next delivery speed shift. And, we've found that they're complementary. The drones will use SSDs as a building block from which to gather selection and launch. And, while Amazon Now will deliver thousands of items within 20 minutes by leveraging micro-fulfillment centers, Prime Air will deliver a much larger selection of items inside a half hour. They will all serve different needs and drive the inflection better as a group than just having one.

Another example of our pursuing multiple paths is Grocery. We started 20 years ago by adding non-perishables (items you'd find in the middle aisles of grocery stores like consumables, canned goods, beauty, etc.). Customers loved the convenience of buying these items online and having them delivered quickly. Unsurprisingly, they asked for broader grocery coverage, with perishables high on their list. We embarked on several bets to try to find the right economic solution for customers and Amazon. We bought Whole Foods Market ("WFM") in 2017, the leader in organic grocery. We've since launched mainstream physical grocery stores (Amazon Fresh), grocery subscriptions for Prime members, and store-within-a-store concepts (with Amazon offering mainstream brands alongside WFM organic brands). Not all of these experiments have worked, but each one has taught us something important.

What's emerged is a clearer picture of what customers want. Our non-perishables grocery business continues to grow quickly. WFM continues to accelerate, with over 550 stores, 100 more coming in the next few years, and an additional smaller format (Daily Shop) serving quick, frequent grocery missions in city neighborhoods that's off to a great start. And, a big breakthrough has been adding perishables into our Same-Day Delivery network. Bringing fresh groceries (e.g. produce, dairy, meat) together with millions of everyday items in a single, fast order is resonating with customers for the value and convenience it provides. Since introducing perishables into Same-Day Delivery in early 2025, perishables sales have grown by over 40 times, and now make up nine of the top ten most-ordered items for Same-Day Delivery where they're available. We have Same-Day fresh food delivery in over 2,300 towns and cities across the country. Our grocery business has grown to over \$150 billion in gross sales in 2025, making Amazon the second-largest grocer in the U.S.

If there's an obvious path to changing your trajectory, take it and run. But, most new jumps forward aren't like that. There's invention and experimentation required, and pursuing multiple paths gives you the best chance to find it.

When you identify disproportionate inflections, bet big. Choosing which inflections are truly seminal versus "just interesting" requires judgment. Reasonable people can disagree. But, if you believe you've found one of these disproportionate shifts, you want to invest as aggressively as you responsibly can. This will create investment spikes that will invite scrutiny, but the game-changers don't typically accommodate smoother investment horizons.

One of these seminal shifts is AI. Every customer experience will be reinvented by AI, and there will be a slew of new experiences only possible because of AI. I've followed the public debate on whether this technology is over-hyped, whether we're in "a bubble," and if the margins and ROIC will be appealing. My strong conviction, at least for Amazon, is that the answers are no, no, and yes. Here are some truths that are hard to debate.

1/ We have never seen a technology more quickly adopted than AI. When ChatGPT launched in November 2022, it reached 100 million users in two months—four times faster than TikTok and 15 times faster than Instagram (ChatGPT already has over 900 million weekly active users). Both OpenAI and Anthropic have revenue run rates reportedly approaching \$30 billion. These are breathtaking numbers for companies this soon after their commercial launches. When Edison opened his first commercial power station in 1882, most people understood it as a better way to light a room. What they couldn't see was that electricity would eventually reorganize every factory, home, and industry on Earth. AI may have comparable impact. The difference is that electricity took 40 years to get where it was going. AI appears to be moving ten times faster.

2/ Amazon is smack in the middle of this land rush, and companies are choosing AWS for AI. Three years after AWS launched commercially, it had a \$58 million revenue run rate. Three years into this AI wave, AWS's AI revenue run rate is over \$15 billion in Q1 2026 (nearly 260 times larger than AWS at that same point)—and ascending rapidly.

Customers are choosing AWS for AI for a few reasons. First, we have broader capabilities than others, with compelling offerings for model-building (SageMaker), high-performance inference with leading selection of frontier models (Bedrock), lower-cost inference (on our custom silicon, Trainium), agent-building (Strands), scalable and secure agent environments (AgentCore), and turnkey agents for coding, software migrations, and most tasks knowledge workers use in their daily routines (Kiro, Transform, and Quick). Second, as customers expand their use of AI, they want their inference to reside near their other applications and data (for latency reasons), and much more of it resides in AWS than anywhere else. Third, as customers expand their AI usage, they consume a lot of additional non-AI services, where AWS also has the broadest and most capable offerings. And fourth, AWS has the strongest security and operational performance of any AI and infrastructure provider. We spend a lot of time listening to customers, and they continue to remark about AWS's advantaged performance as they increasingly move their AI to AWS.

3/ AWS could be growing even faster. AWS added 3.9 gigawatts ("GW") of new power capacity in 2025, expects to double total power capacity by the end of 2027, and is monetizing that capacity as fast as it's installed. In Q4 2025, AWS reported 24% YoY growth with a \$142 billion dollar revenue run rate. That's a lot of absolute growth. And yet, we still have capacity constraints that yield unserved demand. [As an aside, two large AWS customers have already asked if they could buy *all* of our Graviton instance capacity in 2026 (Graviton is our widely-adopted custom CPU chip)—we can't agree to these requests given other customers' needs, but it gives you an idea of the demand.]

4/ Our chips business is on fire, changes the economics for AWS, and will be much larger than most think. Virtually all AI thus far has been done on NVIDIA chips, but a new shift has started. We have a strong partnership with NVIDIA, will always have customers who choose to run NVIDIA, and we will continue to make AWS the best place to run NVIDIA. However, customers want better price-performance. We've seen this movie before. In the CPU space, virtually all of the workloads ran on Intel chips until we invented Graviton in 2018. Graviton, which has up to 40% better price-performance than other x86 processors, is now used expansively by 98% of the top 1,000 EC2 customers. The same story arc is unfolding in AI. Our second version of our custom AI silicon (Trainium2) had about 30% better price-performance than comparable GPUs, and has largely sold out. Trainium3, which just started shipping at the start of 2026 and is 30-40% more price-performant than Trainium2, is nearly fully-subscribed. A significant chunk of Trainium4, which is still about 18 months from broad availability, has already been reserved. And, Amazon Bedrock, AWS's primary (and very fast-growing) inference service, runs most of its inference on Trainium. Demand for Trainium is booming.

Having our own hotly demanded AI chip opens up many possibilities, but perhaps none larger than the ability to lower costs for customers and secure better economics for AWS. At scale, we expect Trainium will save us tens of billions of capex dollars per year, and provide several hundred basis points of operating margin advantage versus relying on others' chips for inference.

Our annual revenue run rate for our chips business (inclusive of Graviton, Trainium, and Nitro—our EC2 NIC) is now over \$20 billion, and growing triple digit percentages YoY. To dimensionalize this versus other chips companies, that run rate is somewhat understated by our currently only monetizing our chips through EC2. If our chips business was a stand-alone business, and sold chips produced this year to AWS

and other third parties (as other leading chips companies do), our annual run rate would be ~\$50 billion. There's so much demand for our chips that it's quite possible we'll sell racks of them to third parties in the future.

5/ The way AWS's cash cycle works is that the faster AWS grows, the more short-term capex we'll spend. AWS has to lay out cash for land, power, buildings, chips, servers, and networking gear in advance of when we can monetize it (typically 6-24 months before we start billing customers, depending on the component). However, these capex investments fund assets with many-year useful lives (30+ years for datacenters; 5-6 years for chips, servers, and networking gear). The FCF and ROIC for these investments are cumulatively quite attractive a couple years after being in service; however, in times of very high growth (like now), where the capex growth meaningfully outpaces the revenue growth, the early-years FCF is challenged until these initial tranches of capacity are being monetized and revenue growth out-paces capex growth. We've been through this cycle with the first big AWS growth wave, and liked the results. We expect to feel similarly about this next wave, with much larger potential downstream revenue and FCF.

6/ We have customer commitments that make our capex investments predictable. We're not investing approximately \$200 billion in capex in 2026 on a hunch. The recent OpenAI commitment (over \$100 billion) is an example of this, but there are several other customer agreements completed (and unannounced), or deep in process. Of the AWS capex we expect to spend in 2026, much of which will be monetized in 2027-2028, we already have customer commitments for a substantial portion of it.

We are willing to make large capex investments and endure short-term FCF headwinds for the substantial medium to long-term FCF surplus. AI is a once-in-a-lifetime opportunity where the current growth is unprecedented and the future growth even bigger. AWS has a significant leadership position with the broadest functionality, strongest security and operational performance, largest share of customers and revenue, strong desire from customers to run their AI in AWS, and an opportunity to build what could be a new pillar for Amazon in chips. We're not going to be conservative in how we play this—we're investing to be the meaningful leader, and our future business, operating income, and FCF will be much larger because of it.

Accept going back to the starting line to redirect the trajectory. When you have a product that's working at scale, one of the hardest decisions to make is to go back to the starting line. It feels like going backwards (because it kind of is). Teams understandably argue that they don't have time to both run the existing service and reimagine everything anew. But, there are reasons to do so, and AI is making it easier and more imperative to go back to the starting line.

For instance, in Amazon Bedrock, as often happens with services that get built rapidly and scale faster than expected, the team realized that it needed a different inference engine than had originally been built. This wasn't a tweak; it required a completely different architecture. The team was operating a large and fast-growing service, so it wasn't ideal. Normally, this sort of activity might take a team of 40 people about a year to carefully build. Instead, the Bedrock team spun up a separable group of six very skilled engineers who were excited about starting over and building on our agentic coding service (Kiro), and delivered this new engine (which we call "Mantle") in 76 days. Mantle has become the backbone of our very successful Bedrock service, which has nearly doubled month-over-month this past March, and processed more tokens in Q1 2026 than all of the tokens processed in the prior years combined.

Alexa is another interesting example. With 600 million active endpoints across devices, cars, offices, Fire TV, and Prime Video, Alexa has scale and a very large customer base. But, when transformative technology like generative AI arrives, and you can build a much more intelligent product than you previously had, you have to pursue it—even if it's disruptive to your team, roadmap, and architecture. The new Alexa+ is so much more capable, useful, and smart than her prior self. But, we had to completely rewire her brain, corresponding intelligence, breadth of knowledge, routing of services and APIs she accesses, and what routines and jobs she could do. The team had to make these changes while also serving the existing large customer base, and execute a roll out that didn't break all the functionality on which so many customers have come to rely. It's been worth it. Customers are talking to Alexa twice as much (and for longer durations across a wider breadth of topics), completing purchases on devices three times more, streaming music 25% more, and using smart home functionality 50% more. Alexa is still early in its journey to be the world's best personal assistant. But, it wouldn't be on its way again without going back to the start.

If you believe that every customer experience will be reinvented in the coming years by AI, it means even the customer experiences that feel most comfy, and are most broadly adopted, will be reimagined. Take our retail business. We believe that customers will always care deeply about massive selection, low prices, very fast delivery, ease of use, and how they're treated. Amazon has built a lot of capabilities that position us well to meet these customer needs for years to come. However, it's not hard to imagine with the emergence of AI, that the interface with which customers want to interact with a retailer could be substantially different over time. The temptation is to just add a little AI to the existing experience. That's a start. But, the trick for leaders, ourselves included, is how to get organized and convicted about going back to the starting line and reimagining your experiences from a clean sheet of paper, assuming you were building with the new technology available. It's easier to say than do. But, it's what we're doing in all of our consumer experiences. It may take us a while to find experiences better than what we have now, and it may take consumers time to adopt these new experiences. But, history shows the "straight line was a lie." Everything gets reinvented. And, if you want to be finding that next zig, you need to be willing to go back to first principles.

Cultivate a culture that can cope with squiggly lines. You may be operating successfully for several years, when an inflection shatters this placidity. Sometimes, teams live in denial that a big change is afoot, as it can be threatening to change course. The historical business landscape is littered with companies that wished away big new inflections to their own detriment (and we have been the beneficiary of a few).

Here's some of what I've learned over the past 28 years at Amazon.

You need to have the right data, mechanisms, and truth-tellers to deeply question what's changed and should therefore be done.

You have to have people and a culture that are comfortable operating with ambiguity as you sort through the new normal.

You need to invent and experiment like crazy. Many of these experiments will fail, and it might feel like you're getting nowhere. But, your culture must possess the tenacity to keep at it.

You need to learn constantly. Each attempt you (and others) make provides insight on what customers really want.

You need to move fast, have teammates that act like true owners, and be scrappy. At Amazon, we talk a lot about operating like the world's biggest startup. It's the primary reason we've worked to flatten our organization the last year, and we're pleased with the improved speed of decision-making and delivery.

You have to be comfortable potentially over-rotating when pursuing one of these big needle-movers. You may end up investing a little more than needed, but it's better to capture most of the transformational upside than to miss much of it by cutting it too fine. And then, if you learn something new that requires course correcting, do it quickly. Inflections aren't usually smooth or calm. They favor the bold and adaptable.

You have to be adept at both inventing the inflections and recognizing when others have started ones that you should pursue as well. We tend to be first-movers, but have had several big successes where we didn't enter first. In each of these cases, we haven't just duplicated what others have done (when we've done that, we haven't succeeded). Rather, we've been inspired by new ideas, and invented unique experiences that appeal to our customers. It takes a different type of grit and resolve than when you're first to an idea.

You have to be willing to reimagine not only every customer experience, but also how you organize and get work done. Challenging conventions that have existed a long time (and worked) is difficult. People sometimes call these "change management" obstacles. That's true, but in times of transformational tumult, they're almost like reexamining your faith.

And, you have to be able to withstand criticism as you're making your way through an inflection. We have historically shared that we're "comfortable being misunderstood for long periods of time." This resolve is required when you're inventing in general, but particularly when you're in the middle of a meaningful change where there's uncertainty and conflicting opinions.

Parting thoughts. Amazon had another strong year in 2025.

Amazon's revenue in 2025 grew 12% year-over-year ("YoY") from \$638 billion to \$717 billion. By segment, North America revenue increased 10% YoY from \$387 billion to \$426 billion, International revenue grew 13% YoY from \$143 billion to \$162 billion, and AWS revenue increased 20% YoY, from \$108 billion to \$129 billion.

Amazon's operating income in 2025 improved 17% YoY, from \$69 billion (an operating margin of 10.8%) to \$80 billion (an operating margin of 11.2%).

Free Cash Flow decreased from \$38 billion to \$11 billion, driven primarily by a year-over-year increase of \$50.7 billion in purchases of property and equipment, net of proceeds from sales and incentives. This FCF change primarily reflects capex investments in artificial intelligence.

It's hard to overstate my optimism for what's ahead.

Our retail business is now approaching \$600 billion in topline, yet roughly 80% of global retail sales still happens in physical stores. That will change.

AWS is at a \$142 billion revenue run rate, and yet 85% of global IT spend remains on-premises. This will change.

Our Advertising offerings continue to grow and deliver strong returns for brands, while newer businesses like Prime Video, Pharmacy, and Grocery are providing unique customer experiences, growing robustly, and improving their economics.

We're on the verge of launching Amazon Leo, just beginning commercial service with Zoox (our autonomous-driving ride-hailing service), and still quite early in what we'll build in robotics.

Across all of this, AI is not a standalone initiative—it's a multiplier. It will reshape every customer experience we offer and unlock entirely new ones. We will build many of these ourselves, and continue making AWS the best place for others to do the same.

Progress will not be linear. There will be moments of acceleration and moments where we adjust course. We will experiment, invest disproportionately behind what matters, and pull back when something isn't working.

This is the environment where Amazon does its best work. We've built a culture that thrives in constantly changing circumstances by being inventive, comfortable operating along multiple paths, revisiting first principles, and committing fully to the ideas that can meaningfully improve customers' lives. As always, we'll optimize for what matters most to customers and shareholders over the long term.

Thank you to our teammates for everything you've done this past year and for what's ahead. I look forward to pursuing this amazing adventure together.

Sincerely,

A handwritten signature in black ink that reads "Andy Jassy". The signature is fluid and cursive, with the first name "Andy" and last name "Jassy" clearly legible.

Andy Jassy
President and Chief Executive Officer
Amazon.com, Inc.

P.S. As we have always done, our original 1997 Shareholder Letter follows. What's written there is as true today as it was in 1997.



1997 LETTER TO SHAREHOLDERS (Reprinted from the 1997 Annual Report)

To our shareholders:

Amazon.com passed many milestones in 1997: by year-end, we had served more than 1.5 million customers, yielding 838% revenue growth to \$147.8 million, and extended our market leadership despite aggressive competitive entry.

But this is Day 1 for the Internet and, if we execute well, for Amazon.com. Today, online commerce saves customers money and precious time. Tomorrow, through personalization, online commerce will accelerate the very process of discovery. Amazon.com uses the Internet to create real value for its customers and, by doing so, hopes to create an enduring franchise, even in established and large markets.

We have a window of opportunity as larger players marshal the resources to pursue the online opportunity and as customers, new to purchasing online, are receptive to forming new relationships. The competitive landscape has continued to evolve at a fast pace. Many large players have moved online with credible offerings and have devoted substantial energy and resources to building awareness, traffic, and sales. Our goal is to move quickly to solidify and extend our current position while we begin to pursue the online commerce opportunities in other areas. We see substantial opportunity in the large markets we are targeting. This strategy is not without risk: it requires serious investment and crisp execution against established franchise leaders.

It's All About the Long Term

We believe that a fundamental measure of our success will be the shareholder value we create over the *long term*. This value will be a direct result of our ability to extend and solidify our current market leadership position. The stronger our market leadership, the more powerful our economic model. Market leadership can translate directly to higher revenue, higher profitability, greater capital velocity, and correspondingly stronger returns on invested capital.

Our decisions have consistently reflected this focus. We first measure ourselves in terms of the metrics most indicative of our market leadership: customer and revenue growth, the degree to which our customers continue to purchase from us on a repeat basis, and the strength of our brand. We have invested and will continue to invest aggressively to expand and leverage our customer base, brand, and infrastructure as we move to establish an enduring franchise.

Because of our emphasis on the long term, we may make decisions and weigh tradeoffs differently than some companies. Accordingly, we want to share with you our fundamental management and decision-making approach so that you, our shareholders, may confirm that it is consistent with your investment philosophy:

- We will continue to focus relentlessly on our customers.
- We will continue to make investment decisions in light of long-term market leadership considerations rather than short-term profitability considerations or short-term Wall Street reactions.
- We will continue to measure our programs and the effectiveness of our investments analytically, to jettison those that do not provide acceptable returns, and to step up our investment in those that work best. We will continue to learn from both our successes and our failures.

- We will make bold rather than timid investment decisions where we see a sufficient probability of gaining market leadership advantages. Some of these investments will pay off, others will not, and we will have learned another valuable lesson in either case.
- When forced to choose between optimizing the appearance of our GAAP accounting and maximizing the present value of future cash flows, we'll take the cash flows.
- We will share our strategic thought processes with you when we make bold choices (to the extent competitive pressures allow), so that you may evaluate for yourselves whether we are making rational long-term leadership investments.
- We will work hard to spend wisely and maintain our lean culture. We understand the importance of continually reinforcing a cost-conscious culture, particularly in a business incurring net losses.
- We will balance our focus on growth with emphasis on long-term profitability and capital management. At this stage, we choose to prioritize growth because we believe that scale is central to achieving the potential of our business model.
- We will continue to focus on hiring and retaining versatile and talented employees, and continue to weight their compensation to stock options rather than cash. We know our success will be largely affected by our ability to attract and retain a motivated employee base, each of whom must think like, and therefore must actually be, an owner.

We aren't so bold as to claim that the above is the "right" investment philosophy, but it's ours, and we would be remiss if we weren't clear in the approach we have taken and will continue to take.

With this foundation, we would like to turn to a review of our business focus, our progress in 1997, and our outlook for the future.

Obsess Over Customers

From the beginning, our focus has been on offering our customers compelling value. We realized that the Web was, and still is, the World Wide Wait. Therefore, we set out to offer customers something they simply could not get any other way, and began serving them with books. We brought them much more selection than was possible in a physical store (our store would now occupy 6 football fields), and presented it in a useful, easy-to-search, and easy-to-browse format in a store open 365 days a year, 24 hours a day. We maintained a dogged focus on improving the shopping experience, and in 1997 substantially enhanced our store. We now offer customers gift certificates, 1-ClickSM shopping, and vastly more reviews, content, browsing options, and recommendation features. We dramatically lowered prices, further increasing customer value. Word of mouth remains the most powerful customer acquisition tool we have, and we are grateful for the trust our customers have placed in us. Repeat purchases and word of mouth have combined to make Amazon.com the market leader in online bookselling.

By many measures, Amazon.com came a long way in 1997:

- Sales grew from \$15.7 million in 1996 to \$147.8 million – an 838% increase.
- Cumulative customer accounts grew from 180,000 to 1,510,000 – a 738% increase.
- The percentage of orders from repeat customers grew from over 46% in the fourth quarter of 1996 to over 58% in the same period in 1997.
- In terms of audience reach, per Media Metrix, our Web site went from a rank of 90th to within the top 20.
- We established long-term relationships with many important strategic partners, including America Online, Yahoo!, Excite, Netscape, GeoCities, AltaVista, @Home, and Prodigy.

Infrastructure

During 1997, we worked hard to expand our business infrastructure to support these greatly increased traffic, sales, and service levels:

- Amazon.com's employee base grew from 158 to 614, and we significantly strengthened our management team.
- Distribution center capacity grew from 50,000 to 285,000 square feet, including a 70% expansion of our Seattle facilities and the launch of our second distribution center in Delaware in November.
- Inventories rose to over 200,000 titles at year-end, enabling us to improve availability for our customers.
- Our cash and investment balances at year-end were \$125 million, thanks to our initial public offering in May 1997 and our \$75 million loan, affording us substantial strategic flexibility.

Our Employees

The past year's success is the product of a talented, smart, hard-working group, and I take great pride in being a part of this team. Setting the bar high in our approach to hiring has been, and will continue to be, the single most important element of Amazon.com's success.

It's not easy to work here (when I interview people I tell them, "You can work long, hard, or smart, but at Amazon.com you can't choose two out of three"), but we are working to build something important, something that matters to our customers, something that we can all tell our grandchildren about. Such things aren't meant to be easy. We are incredibly fortunate to have this group of dedicated employees whose sacrifices and passion build Amazon.com.

Goals for 1998

We are still in the early stages of learning how to bring new value to our customers through Internet commerce and merchandising. Our goal remains to continue to solidify and extend our brand and customer base. This requires sustained investment in systems and infrastructure to support outstanding customer convenience, selection, and service while we grow. We are planning to add music to our product offering, and over time we believe that other products may be prudent investments. We also believe there are significant opportunities to better serve our customers overseas, such as reducing delivery times and better tailoring the customer experience. To be certain, a big part of the challenge for us will lie not in finding new ways to expand our business, but in prioritizing our investments.

We now know vastly more about online commerce than when Amazon.com was founded, but we still have so much to learn. Though we are optimistic, we must remain vigilant and maintain a sense of urgency. The challenges and hurdles we will face to make our long-term vision for Amazon.com a reality are several: aggressive, capable, well-funded competition; considerable growth challenges and execution risk; the risks of product and geographic expansion; and the need for large continuing investments to meet an expanding market opportunity. However, as we've long said, online bookselling, and online commerce in general, should prove to be a very large market, and it's likely that a number of companies will see significant benefit. We feel good about what we've done, and even more excited about what we want to do.

1997 was indeed an incredible year. We at Amazon.com are grateful to our customers for their business and trust, to each other for our hard work, and to our shareholders for their support and encouragement.



Jeffrey P. Bezos
Founder and Chief Executive Officer
Amazon.com, Inc.