



IDEAS ... Michael Lewis

Stock Market Dysfunction Exposed

April 2014

The Wolf Hunters of Wall Street

An Adaptation From 'Flash Boys: A Wall Street Revolt,' by Michael Lewis



From left: Rob Park, Brad Katsuyama and Ronan Ryan. Credit Stefan Ruiz for The New York Times

Before the collapse of the U.S. financial system in 2008, Brad Katsuyama could tell himself that he bore no responsibility for that system. He worked for the Royal Bank of Canada, for a start.

RBC might have been the fifth-biggest bank in North America, by some measures, but it was on nobody's mental map of Wall Street. It was stable and relatively virtuous and soon to be known for having resisted the temptation to make bad subprime loans to Americans or peddle them to ignorant investors. But its management didn't understand just what an afterthought the bank was — on the rare occasions American financiers thought about it at all. Katsuyama's bosses sent him to New York from Toronto in 2002, when he was 23, as part of a "big push" for the bank to become a player on Wall Street. The sad truth was that hardly anyone noticed it. "The people in Canada are always saying, 'We're paying too much for people in the United States,'" Katsuyama says. "What they don't realize is that the reason you have to pay them too much is that no one wants to work for RBC. RBC is a nobody."

Before arriving there as part of the big push, Katsuyama had never laid eyes on Wall Street or New York City. It was his first immersive course in the American way of life, and he was instantly struck by how different it was from the Canadian version. "Everything was to excess," he says. "I met more offensive people in a year than I had in my entire life. People lived beyond their means, and the way they did it was by going into debt. That's what shocked me the most. Debt was a foreign concept in Canada. Debt was evil."

For his first few years on Wall Street, Katsuyama traded U.S. energy stocks and then tech stocks. Eventually he was promoted to run one of RBC's equity-trading groups, consisting of 20 or so traders. The RBC trading floor had a no-jerk rule (though the staff had a more colorful term for it): If someone came in the door looking for a job and sounding like a typical Wall Street jerk, he wouldn't be hired, no matter how much money he said he could make the firm. There was even an expression used to describe the culture: "RBC nice." Although Katsuyama found the expression embarrassingly Canadian, he, too, was RBC nice. The best way to manage people, he thought, was to persuade them that you were good for their careers. He further believed that the only way to get people to believe that you were good for their careers was actually to be good for their careers.

His troubles began at the end of 2006, after RBC paid \$100 million for a U.S. electronic-trading firm called Carlin Financial. In what appeared to Katsuyama to be undue haste, his bosses back in Canada bought Carlin without knowing much about the company or even electronic trading. Now they would receive a crash course. Katsuyama found himself working side by side with a group of American traders who could not have been less suited to RBC's culture. The first day after the merger, Katsuyama got a call from a worried female employee, who whispered, "There is a guy in here with suspenders walking around with a baseball bat in his hands." That turned out to be Carlin's chief executive, Jeremy Frommer, who was, whatever else he was, not RBC nice. Returning to his alma mater, the University at Albany, years later to speak about the secret of his success, Frommer told a group of business students: "It's not just enough to fly in first class; I have to know my friends are flying in coach."

Installed in Carlin's offices, RBC's people in New York were soon gathered to hear a state-of-the-financial-markets address given by Frommer. He stood in front of a flat-panel computer monitor that hung on his wall. "He gets up and says the markets are now all about speed," Katsuyama says. "And then he says, 'I'm going to show you how fast our system is.' He had this guy next to him with a computer keyboard. He said to him, 'Enter an order!' And the guy hit Enter. And the order appeared on the screen so everyone could see it. And Frommer goes: 'See! See how fast that was!!!'" All the guy did was type the name of a stock on a keyboard, and the name was displayed on the screen, the way a letter, once typed, appears on a computer screen.

“Then he goes, ‘Do it again!’ And the guy hits the Enter button on the keyboard again. And everyone nods. It was 5 in the afternoon. The market wasn’t open; nothing was happening. But he was like, ‘Oh, my God, it’s happening in real time!’ ”

Katsuyama couldn’t believe it. He thought: The guy who just sold us our new electronic-trading platform either does not know that his display of technical virtuosity is absurd or, worse, he thinks we don’t know.



A facility in Secaucus, N.J., is home to some IEX servers. Credit Stefan Ruiz for The New York Times

As it happened, at almost exactly the moment Carlin Financial entered Brad Katsuyama’s life, the U.S. stock market began to behave oddly. Before RBC acquired this supposed state-of-the-art electronic-trading firm, Katsuyama’s computers worked as he expected them to. Suddenly they didn’t. It used to be that when his trading screens showed 10,000 shares of Intel offered at \$22 a share, it meant that he could buy 10,000 shares of Intel for \$22 a share. He had only to push a button. By the spring of 2007, however, when he pushed the button to complete a trade, the offers would vanish. In his seven years as a trader, he had always been able to look at the screens

on his desk and see the stock market. Now the market as it appeared on his screens was an illusion.

This made it impossible for Katsuyama to do his job properly. His main role as a trader was to play the middleman between investors who wanted to buy and sell big amounts of stock and the public markets, where the volumes were smaller. Say some investor wanted to sell a block of three million Intel shares, but the markets showed demand for only one million shares: Katsuyama would buy the entire block from the investor, sell off a million shares instantly and then work artfully over the next few hours to unload the other two million. If he didn't know the actual demand in the markets, he couldn't price the larger block. He had been supplying liquidity to the market; now whatever was happening on his screens was reducing his willingness to do that.

By June 2007 the problem had grown too big to ignore. At that point, he did what most people do when they don't understand why their computers aren't working the way they're supposed to: He called tech support. Like tech-support personnel everywhere, their first assumption was that Katsuyama didn't know what he was doing. 'User error' was the thing they'd throw at you," he says. "They just thought of us traders as a bunch of dumb jocks."

Finally he complained so loudly that they sent the developers, the guys who came to RBC in the Carlin acquisition. "They told me it was because I was in New York and the markets were in New Jersey and my market data was slow," Katsuyama says. "Then they said that it was all caused by the fact that there are thousands of people trading in the market. They'd say: 'You aren't the only one trying to do what you're trying to do. There's other events. There's news.'"

If that was the case, he asked them, why did the market in any given stock dry up only when he was trying to trade in it? To make his point, he asked the developers to stand behind him and watch while he traded. "I'd say: 'Watch closely. I am about to buy 100,000 shares of AMD. I am willing to pay \$15 a share. There are currently 100,000 shares of AMD being offered at \$15 a share — 10,000 on BATS, 35,000 on the New York Stock Exchange, 30,000 on Nasdaq and 25,000 on Direct Edge.' You could see it all on the screens. We'd all sit there and stare at the screen, and I'd have my finger over the Enter button. I'd count out loud to five. . . .

"One. . . .

"Two. . . . See, nothing's happened.


"Three. . . . Offers are still there at 15. . . .

"Four. . . . Still no movement. . . .

"Five.' Then I'd hit the Enter button, and — boom! — all hell would break loose. The offerings would all disappear, and the stock would pop higher."

At which point he turned to the developers behind him and said: "You see, I'm the event. I am the news."

High-frequency-trading activity is not constant; it occurs in microbursts. The line at the bottom of this graphic is the stock-market activity involving General Electric shares over 100 milliseconds (one-tenth of a second) at 12:44 p.m. on Dec. 19, 2013. The gray box magnifies a five-millisecond window, during which GE experienced heavy bid and offer activity and a total of 44 trades.



While You Were Blinking ... Credit Graphic: CLEVER^oFRANKE. Data source: IEX.

To that, they had no response. Katsuyama suspected the culprit was Carlin's setup. "As the market problem got worse," he says, "I started to just assume my real problem was with how bad their technology was."

But as he talked to Wall Street investors, he came to realize that they were dealing with the same problem. He had a good friend who traded stocks at a big-time hedge fund in Stamford, Conn., called SAC Capital, which was famous (and soon to be infamous) for being one step ahead of the U.S. stock market. If anyone was going to know something about the market that Katsuyama didn't know, he figured, it would be someone there. One spring morning, he took the train up to Stamford and spent the day watching his friend trade. Right away he saw that, even though his friend was using software supplied to him by Goldman Sachs and Morgan Stanley and the other big firms, he was experiencing exactly the same problem as RBC: He would hit a button to buy or sell a stock, and the market would move away from him. "When I see this guy trading, and he was getting screwed — I now see that it isn't just me. My frustration is the market's frustration. And I was like, 'Whoa, this is serious.'"

<p>

People always assumed that because he was Asian, Brad Katsuyama must be a computer wizard. In reality, he couldn't (or wouldn't) even program his own DVR. What he had was an ability to distinguish between computer people who actually knew what they were talking about and those who didn't. So he wasn't exactly shocked when RBC finally gave up looking for someone to run its mess of an electronic-trading operation and asked him if he would take over and try to fix it. He shocked his friends and colleagues, however, when he agreed to do it, because A) he had a safe and cushy \$1.5-million-a-year job running the human traders, and B) RBC had nothing to add to electronic trading. The market was cluttered; big investors had use for only so many trading algorithms sold by brokers; and Goldman Sachs and Morgan Stanley and Credit Suisse had long since overrun that space.

So Katsuyama was in charge of a business called electronic trading — with only Carlin's inferior software to sell. What he had, instead, was a fast-growing pile of unanswered questions. Between the public stock exchanges and the dark pools — private exchanges created by banks and brokers that did not have to report in real time what trading activities took place within them — why were there now nearly 60 different places, most of them in New Jersey, where you could buy any listed stock? Why did one public exchange pay you to do something — sell shares, say — when another exchange charged you to do the same exact thing? Why was the market displayed on Wall Street trading screens an illusion?

He hired Rob Park, a gifted technologist, to explain to him what actually happened inside all these new Wall Street black boxes, and together they set out to assemble a team to investigate the U.S. stock market. Once he had a team in place, Katsuyama persuaded his superiors at RBC to conduct what amounted to a series of experiments. For the next several months, he and his people would trade stocks not to make money but to test theories. RBC agreed to let his team lose up to \$10,000 a day to figure out why the market in any given stock vanished the moment RBC tried to trade in it. Katsuyama asked Park to come up with some theories.

They started with the public markets — 13 stock exchanges scattered over four different sites run by the New York Stock Exchange, Nasdaq, BATS and Direct Edge. Park's first theory was that the exchanges weren't simply bundling all the orders at a given price but arranging them in some kind of sequence. You and I might each submit an order to buy 1,000 shares of Intel at \$30 a share, but you might somehow obtain the right to cancel your order if my order was filled. "We started getting the idea that people were canceling orders," Park says. "That they were just phantom orders."

Katsuyama tried sending orders to a single exchange, fairly certain that this would prove that some, or maybe even all, of the exchanges were allowing these phantom orders. But no: To his surprise, an order sent to a single exchange enabled him to buy everything on offer. The market as it appeared on his screens was, once again, the market. "I thought, [expletive], there goes that theory," Katsuyama says. "And that's our only theory."

It made no sense: Why would the market on the screens be real if you sent your order to only one exchange but prove illusory when you sent your order to all the exchanges at once? The team began to send orders into various combinations of exchanges. First the New York Stock Exchange and Nasdaq. Then N.Y.S.E. and Nasdaq and BATS. Then N.Y.S.E., Nasdaq BX, Nasdaq and BATS. And so on. What came back was a further mystery. As they increased the number of exchanges, the percentage of the order that was filled decreased; the more places they tried to buy stock from, the less stock they were actually able to buy. "There was one exception,"

Katsuyama says. “No matter how many exchanges we sent an order to, we always got 100 percent of what was offered on BATS.” Park had no explanation, he says. “I just thought, BATS is a great exchange!”



This box kept at the facility in Secaucus, N.J., contains a 38-mile coil of fiber-optic cable that creates a slight delay in the processing of orders, which levels the playing field among traders. Credit Stefan Ruiz for The New York Times

One morning, while taking a shower, Rob Park came up with another theory. He was picturing a bar chart he had seen that showed the time it took orders to travel from Brad Katsuyama’s trading desk in the World Financial Center to the various exchanges.

The increments of time involved were absurdly small: In theory, the fastest travel time, from Katsuyama’s desk in Manhattan to the BATS exchange in Weehawken, N.J., was about two milliseconds, and the slowest, from Katsuyama’s desk to the Nasdaq exchange in Carteret, N.J., was around four milliseconds. In practice, the times could vary much more than that, depending on network traffic, static and glitches in the equipment between any two points. It takes 100 milliseconds to blink quickly — it was hard to believe that a fraction of a blink of an eye could have any real market consequences.

Trading in the Fast Lane



In the microseconds it takes a high-frequency trader — depicted in blue — to reach the various stock exchanges housed in these New Jersey towns, the conventional trader's order, theoretically, makes it only as far as the red line. The time differences — now under investigation by New York's attorney general — can be financially advantageous in a number of ways. Credit Graphic: CLEVER°FRANKE. Data source: IEX.

Allen Zhang, whom Katsuyama and Park viewed as their most talented programmer, wrote a program that built delays into the orders Katsuyama sent to exchanges that were faster to get to, so that they arrived at exactly the same time as they did at the exchanges that were slower to get

to. “It was counterintuitive,” Park says, “because everyone was telling us it was all about faster. We had to go faster, and we were slowing it down.” One morning they sat down to test the program. Ordinarily when you hit the button to buy but failed to get the stock, the screens lit up red; when you got only some of the stock you were after, the screens lit up brown; and when you got everything you asked for, the screens lit up green.

The screens lit up green.

“It’s 2009,” Katsuyama says. “This had been happening to me for almost two years. There’s no way I’m the first guy to have figured this out. So what happened to everyone else?” The question seemed to answer itself: Anyone who understood the problem was making money off it.

Now he and RBC had a tool to sell to investors: The program Zhang wrote to build delays into the stock-exchange orders. The tool enabled traders like Katsuyama to do the job they were meant to do — take risk on behalf of the big investors who wanted to trade big chunks of stock. They could once again trust the market on their screens. The tool needed a name. The team stewed over this, until one day a trader stood up at his desk and hollered: “Dude, you should just call it Thor! The hammer!” Someone was assigned to figure out what Thor might be an acronym for, and some words were assembled, but no one remembered them. The tool was always just Thor. “I knew we were onto something when Thor became a verb,” Katsuyama says. “When I heard guys shouting, ‘Thor it!’ ”

The other way he knew they were on to something was from conversations he had with a few of the world’s biggest money managers. The first visit Katsuyama and Park made was to Mike Gitlin, who oversaw global trading for billions of dollars in assets for the money-management firm T. Rowe Price. The story they told didn’t come to Gitlin as a complete shock. “You could see that something had just changed,” Gitlin says. “You could see that when you were trading a stock, the market knew what you were going to do, and it was going to move against you.” But what Katsuyama described was a far more detailed picture of the market than Gitlin had ever considered — and in that market, all the incentives were screwy. The Wall Street brokerage firm that was deciding where to send T. Rowe Price’s buy and sell orders had a great deal of power over how and where those orders were submitted. Some exchanges paid brokerages for their orders; others charged for those orders. Did that influence where the broker decided to send an order, even when it didn’t sync with the interests of the investors the broker was supposed to represent? No one could say. Another wacky incentive was “payment for order flow.” As of 2010, every American brokerage and all the online brokers effectively auctioned their customers’ stock-market orders. The online broker TD Ameritrade, for example, was paid hundreds of millions of dollars each year to send its orders to a hedge fund called Citadel, which executed the orders on behalf of TD Ameritrade. Why was Citadel willing to pay so much to see the flow? No one could say with certainty what Citadel’s advantage was.

Katsuyama and his team did measure how much more cheaply they bought stock when they removed the ability of some other unknown trader to front-run them. For instance, they bought 10 million shares of Citigroup, then trading at roughly \$4 per share, and saved \$29,000 — or less than 0.1 percent of the total price. “That was the invisible tax,” Park says. It sounded small until you realized that the average daily volume in the U.S. stock market was \$225 billion. The same tax rate applied to that sum came to nearly \$160 million a day. “It was so insidious because you couldn’t see it,” Katsuyama says. “It happens on such a granular level that even if you tried to

line it up and figure it out, you wouldn't be able to do it. People are getting screwed because they can't imagine a microsecond."

Ronan Ryan didn't look like a Wall Street trader. He had pale skin and narrow, stooped shoulders and the uneasy caution of a man who has survived one potato famine and is expecting another. He also lacked the Wall Street trader's ability to seem more important and knowledgeable than he actually was. Yet from the time he first glimpsed a Wall Street trading floor, in his early 20s, Ryan badly wanted to work there. "It's hard not to get enamored of being one of these Wall Street guys who people are scared of and make all this money," he says.

Born and raised in Dublin, he moved to America in 1990, when he was 16. Six years later, his father was recalled to Ireland; Ronan stayed behind. He didn't think of Ireland as a place anyone would ever go back to if given the choice, and he embraced his version of the American dream. After graduating from Fairfield University in 1996, he sent letters to all the Wall Street banks, but he received just one flicker of interest from what, even to his untrained eyes, was a vaguely criminal, pump-and-dump penny-stock brokerage firm.

Eventually he met another Irish guy who worked in the New York office of MCI Communications, the big telecom company. "He gave me a job strictly because I was Irish," Ryan says.



Photo ... Brad Katsuyama and his team at IEX. Credit Stefan Ruiz for The New York Times

He had always been handy, but he never actually studied anything practical. He knew next to nothing about technology. Now he started to learn all about it. “It’s pretty captivating, when you take the nerdiness out of it” and figure out how stuff works, he says. How a copper circuit conveyed information, compared with a glass fiber. How a switch made by Cisco compared with a switch made by Juniper. Which hardware companies made the fastest computer equipment and which buildings in which cities contained floors that could withstand the weight of that equipment (old manufacturing buildings were best). He also learned how information actually traveled from one place to another — not in a straight line run by a single telecom carrier, usually, but in a convoluted path run by several. “When you make a call to New York from Florida, you have no idea how many pieces of equipment you have to go through for that call to happen. You probably just think it’s like two cans and a piece of string. But it’s not.” A circuit that connected New York City to Florida would have Verizon on the New York end, AT&T on the Florida end and MCI in the middle; it would zigzag from population center to population center.

Ryan hadn’t been able to find a job on Wall Street, but by 2005 his clients were more likely than ever to be big Wall Street banks. He spent entire weeks inside Goldman Sachs and Lehman Brothers and Deutsche Bank, finding the best routes for their fiber and the best machines to execute their stock-market trades.

In 2005, he went to work for BT Radianz, a company that was born of 9/11, after the attacks on the World Trade Center knocked out big pieces of Wall Street’s communication system. The company promised to build a system less vulnerable to outside attack. Ryan’s job was to sell the financial world on the idea of subcontracting its information networks to Radianz. In particular, he was meant to sell the banks on “co-locating” their computers in Radianz’s data center in Nutley, N.J., to be closer, physically, to where the stock exchanges were located.

‘Thank God,’ a big investor said, ‘finally there’s someone who knows something about high-frequency trading who isn’t an Area 51 guy.’

Not long after he started his job at Radianz, Ryan received an inquiry from a hedge fund based in Kansas City, Kan. The caller said he worked at a stock-market trading firm called Bountiful Trust and that he heard Ryan was an expert at moving financial data from one place to another. Bountiful Trust had a problem: In making trades between Kansas City and New York, it took too long to determine what happened to the firm’s orders — that is, what stocks had been bought and sold. They also noticed that, increasingly, when they placed their orders, the market was vanishing on them. “He says, ‘My latency time is 43 milliseconds,’ ” Ryan recalls. “And I said, ‘What the hell is a millisecond?’ ”

“Latency” was simply the time between the moment a signal was sent and when it was received. Several factors determined the latency of a trading system: the boxes, the logic and the lines. The boxes were the machinery the signals passed through on their way from Point A to Point B: the computer servers and signal amplifiers and switches. The logic was the software, the code instructions that operated the boxes. Ryan didn’t know much about software, except that more and more it seemed to be written by guys with thick Russian accents. The lines were the glass fiber-optic cables that carried the information from one box to another. The single-biggest determinant of speed was the length of the fiber, or the distance the signal needed to travel. Ryan didn’t know what a millisecond was, but he understood the problem with this Kansas City hedge fund: It was in Kansas City. Light in a vacuum travels at 186,000 miles per second or, put another way, 186 miles a millisecond. Light inside fiber bounces off the walls and travels at only

about two-thirds of its theoretical speed. “Physics is physics — this is what the traders didn’t understand,” Ryan says.

By the end of 2007, Ryan was making hundreds of thousands of dollars a year building systems to make stock-market trades faster. He was struck, over and over again, by how little those he helped understood the technology they were using. Beyond that, he didn’t even really know much about his clients. The big banks — Goldman Sachs, Citigroup — everyone had heard of. Others — Citadel, Getco — were famous on a small scale. He learned that some of these firms were hedge funds, which meant they took money from outside investors. But most of them were proprietary firms, or prop shops, trading only their own founders’ money. A huge number of the outfits he dealt with — Hudson River Trading, Eagle Seven, Simplex Investments, Evolution Financial Technologies, Cooperfund, DRW — no one had ever heard of, and the firms obviously intended to keep it that way. The prop shops were especially strange, because they were both transient and prosperous. “They’d be just five guys in a room. All of them geeks. The leader of each five-man pack is just an arrogant version of that geek.” One day a prop shop was trading; the next, it closed, and all the people in it moved on to work for some big Wall Street bank. One group of guys Ryan saw over and over: four Russian, one Chinese. The arrogant Russian guy, clearly the leader, was named Vladimir, and he and his boys bounced from prop shop to big bank and back to prop shop, writing the computer code that made the actual stock-market trading decisions, which made high-frequency trading possible. Ryan watched them meet with one of the most senior guys at a big Wall Street bank that hoped to employ them — and the Wall Street big shot sucked up to them. “He walks into the meeting and says, ‘I’m always the most important man in the room, but in this case, Vladimir is.’”

“I needed someone from the industry to verify that what I was saying was real,” Katsuyama says. He needed, specifically, someone from deep inside the world of high-frequency trading. He spent the better part of a year cold-calling strangers in search of a high-frequency-trading strategist willing to defect. He now suspected that every human being who knew how high-frequency traders made money was making too much money doing it to stop and explain what they were doing. He needed to find another way in.

In the fall of 2009, Katsuyama’s friend at Deutsche Bank mentioned this Irish guy who seemed to be the world’s expert at helping the world’s fastest stock-market traders be faster. Katsuyama called Ronan Ryan and invited him to interview for a job on the RBC trading floor. In his interview, Ryan described what he witnessed inside the exchanges: The frantic competition for nanoseconds, clients’ trying to get their machines closer to the servers within the exchanges, the tens of millions being spent by high-frequency traders for tiny increments of speed. The U.S. stock market was now a class system of haves and have-nots, only what was had was not money but speed (which led to money). The haves paid for nanoseconds; the have-nots had no idea that a nanosecond had value. The haves enjoyed a perfect view of the market; the have-nots never saw the market at all. “I learned more from talking to him in an hour than I learned from six months of reading about [high-frequency trading],” Katsuyama says. “The second I met him, I wanted to hire him.”

He wanted to hire him without being able to fully explain, to his bosses or even to Ryan, what he wanted to hire him for. He couldn’t very well call him vice president in charge of explaining to my clueless superiors why high-frequency trading is a travesty. So he called him a high-frequency-trading strategist. And Ryan finally landed his job on a Wall Street trading floor.

Katsuyama and his team were having trouble turning Thor into a product RBC could sell to investors. They had no control over the path the signals took to get to the exchanges or how much traffic was on the network. Sometimes it took four milliseconds for their orders to arrive at the New York Stock Exchange; other times, it took seven milliseconds. In short, Thor was inconsistent — because, Ryan explained, the paths the electronic signals took from Katsuyama’s desk to the various exchanges were inconsistent. The signal sent from Katsuyama’s desk arrived at the New Jersey exchanges at different times because some exchanges were farther from Katsuyama’s desk than others. The fastest any high-speed trader’s signal could travel from the first exchange it reached to the last one was 465 microseconds, or one two-hundredth of the time it takes to blink. (A microsecond is a millionth of a second.) That is, for Katsuyama’s trading orders to interact with the market as displayed on his trading screens, they needed to arrive at all the exchanges within a 465-microsecond window.



Photo ... The IEX office in Manhattan. Credit Stefan Ruiz for The New York Times

To make his point, Ryan brought in oversize maps of New Jersey showing the fiber-optic networks built by telecom companies. The maps told a story: Any trading signal that originated in Lower Manhattan traveled up the West Side Highway and out the Lincoln Tunnel. Perched immediately outside the tunnel, in Weehawken, N.J., was the BATS exchange. From BATS the routes became more complicated, as they had to find their way through the clutter of the Jersey suburbs. “New Jersey is now carved up like a Thanksgiving turkey,” Ryan says. One way or

another, they traveled west to Secaucus, the location of the Direct Edge family of exchanges owned in part by Goldman Sachs and Citadel, and south to the Nasdaq family of exchanges in Carteret. The New York Stock Exchange, less than a mile from Katsuyama's desk, appeared to be the stock market closest to him — but Ryan's maps showed the incredible indirection of fiber-optic cable in Manhattan. "To get from Liberty Plaza to 55 Water Street, you might go through Brooklyn," he explained. "You can go 50 miles to get from Midtown to Downtown. To get from a building to a building across the street, you could travel 15 miles."

To Katsuyama the maps explained, among other things, why the market on BATS had proved so accurate. The reason they were always able to buy or sell 100 percent of the shares listed on BATS was that BATS was always the first stock market to receive their orders. News of their buying and selling hadn't had time to spread throughout the marketplace. Inside BATS, high-frequency-trading firms were waiting for news that they could use to trade on the other exchanges. BATS, unsurprisingly, had been created by high-frequency traders.

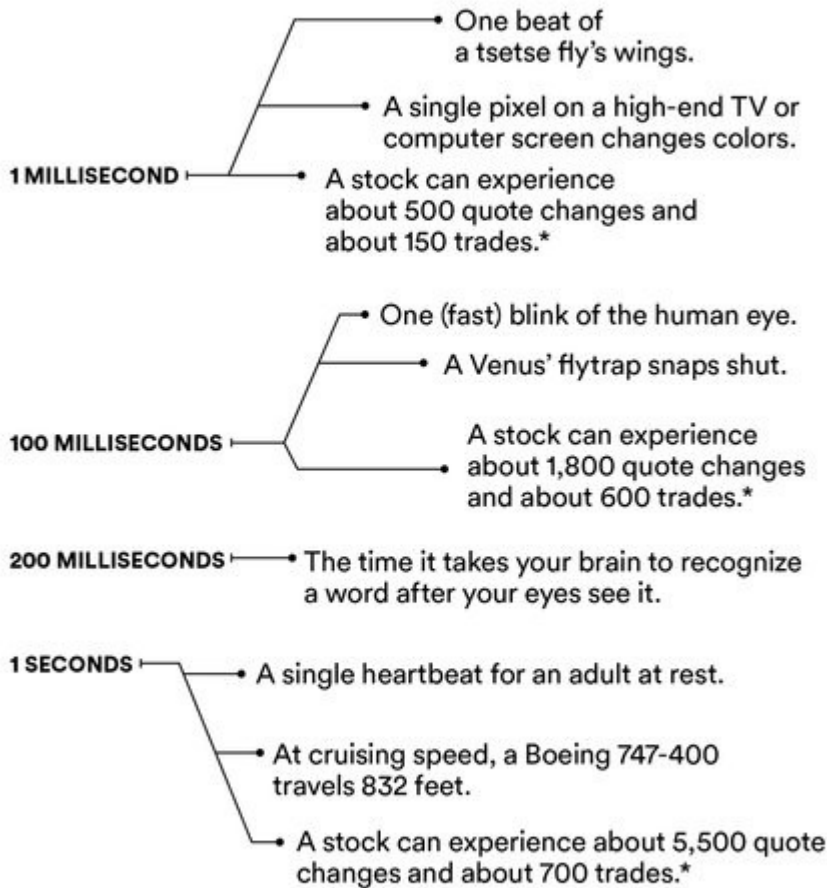
Eventually Brad Katsuyama came to realize that the most sophisticated investors didn't know what was going on in their own market. Not the big mutual funds, Fidelity and Vanguard. Not the big money-management firms like T. Rowe Price and Capital Group. Not even the most sophisticated hedge funds. The legendary investor David Einhorn, for instance, was shocked; so was Dan Loeb, another prominent hedge-fund manager. Bill Ackman runs a famous hedge fund, Pershing Square, that often buys large chunks of companies. In the two years before Katsuyama turned up in his office to explain what was happening, Ackman had started to suspect that people might be using the information about his trades to trade ahead of him. "I felt that there was a leak every time," Ackman says. "I thought maybe it was the prime broker. It wasn't the kind of leak that I thought." A salesman at RBC who marketed Thor recalls one big investor calling to say, "You know, I thought I knew what I did for a living, but apparently not, because I had no idea this was going on."

Katsuyama and Ryan between them met with roughly 500 professional stock-market investors who controlled many trillions of dollars in assets. Most of them had the same reaction: They knew something was very wrong, but they didn't know what, and now that they knew, they were outraged. Vincent Daniel, a partner at Seawolf, took a long look at this unlikely pair — an Asian-Canadian guy from a bank no one cared about and an Irish guy who was doing a fair impression of a Dublin handyman — who just told him the most incredible true story he had ever heard and said, "Your biggest competitive advantage is that you don't want to [expletive] me."

Trust on Wall Street was still — just — possible. The big investors who trusted Katsuyama began to share whatever information they could get their hands on from their other brokers. For instance, several demanded to know from their other Wall Street brokers what percentage of the trades executed on their behalf were executed inside the brokers' dark pools. Goldman Sachs and Credit Suisse ran the most prominent dark pools, but every brokerage firm strongly encouraged investors who wanted to buy or sell big chunks of stock to do so in that firm's dark pool. In theory, the brokers were meant to find the best price for their customers. If the customer wanted to buy shares in Chevron, and the best price happened to be on the New York Stock Exchange, the broker was not supposed to stick the customer with a worse price inside its own dark pool. But the dark pools were opaque. Their rules were not published. No outsider could see what went on inside them. It was entirely possible that a broker's own traders were trading against the customers in its dark pool: There were no rules against doing that. And while the brokers often protested that there were no conflicts of interest inside their dark pools, all the dark pools

exhibited the same strange property: A huge percentage of the customer orders sent into a dark pool were executed inside the pool. Even giant investors simply had to take it on faith that Goldman Sachs or Merrill Lynch acted in their interests, despite the obvious financial incentives not to do so. As Mike Gitlin of T. Rowe Price says: "It's just very hard to prove that any broker dealer is routing the trades to someplace other than the place that is best for you. You couldn't see what any given broker was doing." If an investor as large as T. Rowe Price, which acted on behalf of millions of investors, had trouble obtaining the information it needed to determine if its brokers had acted in their interest, what chance did the little guy have?

WHEN EVERY MILLISECOND COUNTS



*These are rounded numbers for actual stock symbols at their most-active period on Dec. 19, 2013, according to IEX.

The deep problem with the system was a kind of moral inertia. So long as it served the narrow self-interests of everyone inside it, no one on the inside would ever seek to change it, no matter

how corrupt or sinister it became — though even to use words like “corrupt” and “sinister” made serious people uncomfortable, so Katsuyama avoided them. Maybe his biggest concern, when he spoke to investors, was that he’d be seen as just another nut with a conspiracy theory. One compliment that made him happiest was when a big investor said, “Thank God, finally there’s someone who knows something about high-frequency trading who isn’t an Area 51 guy.” It took him a while to figure out that fate and circumstance had created for him a dramatic role, which he was obliged to play. One night he turned to his wife, Ashley, and said: “It feels like I’m an expert in something that badly needs to be changed. I think there’s only a few people in the world who can do anything about this. If I don’t do something right now — me, Brad Katsuyama — there’s no one to call.”

In May 2011, the small team Katsuyama created — Ronan Ryan, Rob Park and a couple of others — sat around a table in his office, surrounded by the applications of past winners of The Wall Street Journal’s Technology Innovation Awards. RBC’s marketing department had informed them of the awards the day before submissions were due and suggested they put themselves forward — so they were scrambling to figure out which of several categories they belonged in and how to make Thor sound life-changing. “There were papers everywhere,” Park says. “No one sounded like us. There were people who had, like, cured cancer.”

“It was stupid,” Katsuyama says, “there wasn’t even a category to put us into. I think we ended up applying under Other.”

With the purposelessness of the exercise hanging in the air, Park said, “I just had a sick idea.” His idea was to license the technology to one of the exchanges. The line between Wall Street brokers and exchanges had blurred. For a few years, the big Wall Street banks had been running their own private exchanges. The stock exchanges, for their part, were making a bid (that ultimately failed) to become brokers. The bigger ones now offered a service that enabled brokers to simply hand them their stock-market orders, which they would then route — to their own exchange, of course, but also to other exchanges. The service was used mainly by small regional brokerage firms that didn’t have their own routers, but this brokeragelike service opened up, at least in Park’s mind, a new possibility. If just one exchange was handed the tool for protecting investors from market predators, the small brokers from around the country might flock to it, and it might become the mother of all exchanges.

Forget that, Katsuyama said. “Let’s just create our own stock exchange.”

“We just sat there for a while,” Park says, “kind of staring at each other. Create your own stock exchange. What does that even mean?”

A few weeks later Katsuyama flew to Canada and tried to sell his bosses on the idea of an RBC-led stock exchange. Then in the fall of 2011, he canvassed a handful of the world’s biggest money managers (Capital Group, T. Rowe Price, BlackRock, Wellington, Southeastern Asset Management) and some of the most influential hedge funds (run by David Einhorn, Bill Ackman, Daniel Loeb). They all had the same reaction. They loved the idea of a stock exchange that protected investors from Wall Street’s predators. They also thought that for a new stock exchange to be credibly independent of Wall Street, it could not be created by a Wall Street bank. Not even a bank as nice as RBC. If Katsuyama wanted to create the mother of all stock exchanges, he would need to quit his job and do it on his own.

The challenges were obvious. He would need to find money. He would need to persuade a lot of highly paid people to quit their Wall Street jobs to work for tiny fractions of their current salaries — and possibly even supply the capital to pay themselves to work. “I was asking: Can I get the people I need? How long can we survive without getting paid? Will our significant others let us do this?” They did, and Katsuyama’s team followed him to his new venture.

But he also needed to find out if the nine big Wall Street banks that controlled nearly 70 percent of all investor orders would be willing to send those orders to a truly safe exchange. It would be far more difficult to start an exchange premised on fairness if the banks that controlled a vast majority of the customers’ orders were not committed to fairness themselves.

Back in 2008, when it first occurred to Brad Katsuyama that the stock market had become a black box whose inner workings eluded ordinary human understanding, he went looking for technologically gifted people who might help him open the box and understand its contents. He’d started with Rob Park and Ronan Ryan, then others. One was a 20-year-old Stanford junior named Dan Aisen, whose résumé Katsuyama discovered in a pile at RBC. The line that leapt out at him was “Winner of Microsoft’s College Puzzle Challenge.” Every year, Microsoft sponsored this one-day, 10-hour national brain-twisting marathon. It attracted more than a thousand young math and computer-science types. Aisen and three friends competed in 2007 and won the whole thing. “It’s kind of a mix of cryptography, ciphers and Sudoku,” Aisen says. To be really good at it, a person needed not only technical skill but also exceptional pattern recognition. “There’s some element of mechanical work and some element of ‘Aha!’ ” Aisen says. Katsuyama gave Aisen both a job and a nickname, the Puzzle Master, soon shortened, by RBC’s traders, to Puz. Puz was one of the people who had helped create Thor.

‘We just sat there for a while, kind of staring at each other. Create your own stock exchange. What does that even mean?’

Puz’s peculiar ability to solve puzzles was suddenly even more relevant. Creating a new stock exchange is a bit like creating a casino: Its creator needs to ensure that the casino cannot in some way be exploited by the patrons. Or at worst, he needs to know exactly how his system might be gamed, so that he might monitor the exploitation — as a casino follows card counters at the blackjack tables. “You are designing a system,” Puz says, “and you don’t want the system to be gameable.” The trouble with the stock market — with all of the public and private exchanges — was that they were fantastically gameable, and had been gamed: first by clever guys in small shops, and then by prop traders who moved inside the big Wall Street banks. That was the problem, Puz thought. From the point of view of the most sophisticated traders, the stock market wasn’t a mechanism for channeling capital to productive enterprise but a puzzle to be solved. “Investing shouldn’t be about gaming a system,” he says. “It should be about something else.”

The simplest way to design a stock exchange that could not be gamed was to hire the very people best able to game it and encourage them to take their best shots. Katsuyama didn’t know any other national puzzle champions, but Puz did. The only problem was that none of them had ever worked inside a stock exchange. “The Puzzle Masters needed a guide,” Katsuyama says.

Enter Constantine Sokoloff, who had helped build Nasdaq’s matching engine — the computer that matched buyers and sellers. Sokoloff was Russian, born and raised in a city on the Volga River. He had an explanation for why so many of his countrymen wound up in high-frequency trading. The old Soviet educational system channeled people into math and science. And the Soviet-controlled economy was horrible and complicated but riddled with loopholes, an

environment that left those who mastered it well prepared for Wall Street in the early 21st century. “We had this system for 70 years,” Sokoloff says. “The more you cultivate a class of people who know how to work around the system, the more people you will have who know how to do it well.”

The Puzzle Masters might not have thought of it this way at first, but in trying to design their exchange so that investors who came to it would remain safe from predatory traders, they were also divining the ways in which high-frequency traders stalked their prey. For example, these traders had helped the public stock exchanges create all sorts of complicated “order types.” The New York Stock Exchange, for one, had created an order type that traded only if the order on the other side of it was smaller than itself — the purpose of which seemed to be to protect high-frequency traders from buying a small number of shares from an investor who was about to depress the market in these shares with a huge sale.

As they worked through the order types, the Puzzle Masters created a taxonomy of predatory behavior in the stock market. Broadly speaking, it appeared as if there were three activities that led to a vast amount of grotesquely unfair trading. The first they called electronic front-running — seeing an investor trying to do something in one place and racing ahead of him to the next (what had happened to Katsuyama when he traded at RBC). The second they called rebate arbitrage — using the new complexity to game the seizing of whatever legal kickbacks, called rebates within the industry, the exchange offered without actually providing the liquidity that the rebate was presumably meant to entice. The third, and probably by far the most widespread, they called slow-market arbitrage. This occurred when a high-frequency trader was able to see the price of a stock change on one exchange and pick off orders sitting on other exchanges before those exchanges were able to react. This happened all day, every day, and very likely generated more billions of dollars a year than the other strategies combined.

All three predatory strategies depended on speed. It was Katsuyama who had the crude first idea to counter them: Everyone was fighting to get in as close to the exchange as possible — why not push them as far away as possible? Put ourselves at a distance, but don’t let anyone else be there. The idea was to locate their exchange’s matching engine at some meaningful distance from the place traders connected to the exchange (called the point of presence) and to require anyone who wanted to trade to connect to the exchange at that point of presence. If you placed every participant in the market far enough away from the exchange, you could eliminate most, and maybe all, of the advantages created by speed. Their matching engine, they already knew, would be located in Weehawken (where they’d been offered cheap space in a data center). The only question was: Where to put the point of presence? “Let’s put it in Nebraska,” someone said, but they all knew it would be harder to get the already reluctant Wall Street banks to connect to their market if the banks had to send people to Omaha to do it. Actually, though, it wasn’t necessary for anyone to move to Nebraska. The delay needed only to be long enough for their new exchange, once it executed some part of a customer’s buy order, to beat high-frequency traders in a race to the shares at any other exchange — that is, to prevent electronic front-running. The necessary delay turned out to be 320 microseconds; that was the time it took them, in the worst case, to send a signal to the exchange farthest from them, the New York Stock Exchange in Mahwah. Just to be sure, they rounded it up to 350 microseconds.

To create the 350-microsecond delay, they needed to keep the new exchange roughly 38 miles from the place the brokers were allowed to connect to the exchange. That was a problem. Having cut one very good deal to put the exchange in Weehawken, they were offered another: to

establish the point of presence in a data center in Secaucus. The two data centers were less than 10 miles apart and already populated by other stock exchanges and all the high-frequency stock-market traders. (“We’re going into the lion’s den,” Ryan said.) A bright idea came from a new employee, James Cape, who had just joined them from a high-frequency-trading firm: Coil the fiber. Instead of running straight fiber between the two places, why not coil 38 miles of fiber and stick it in a compartment the size of a shoe box to simulate the effects of the distance. And that’s what they did.

Creating fairness was remarkably simple. They would not sell to any one trader or investor the right to put his computers next to the exchange or special access to data from the exchange. They would pay no rebates to brokers or banks that sent orders; instead, they would charge both sides of any trade the same amount: nine one-hundredths of a cent per share (known as nine mils). They’d allow just three order types: market, limit and Mid-Point Peg, which meant that the investor’s order rested in between the current bid and offer of any stock. If the shares of Procter & Gamble were quoted in the wider market at 80–80.02 (you can buy at \$80.02 or sell at \$80), a Mid-Point Peg order would trade only at \$80.01. “It’s kind of like the fair price,” Katsuyama says.

Finally, to ensure that their own incentives remained as closely aligned as they could be with those of investors, the new exchange did not allow anyone who could trade directly on it to own any piece of it: Its owners were all ordinary investors who needed first to hand their orders to brokers.

But the big Wall Street banks that controlled a majority of all stock-market investor orders played a more complicated role than an online broker like TD Ameritrade. The Wall Street banks controlled not only the orders, and the informational value of those orders, but also dark pools in which those orders might be executed. The banks took different approaches to milking the value of their customers’ orders. All of them tended to send the orders first to their own dark pools before routing them out to the wider market. Inside the dark pool, the bank could trade against the orders itself; or it could sell special access to the dark pool to high-frequency traders. Either way, the value of the customers’ orders was monetized — by the big Wall Street bank for the big Wall Street bank. If the bank was unable to execute an order in its own dark pool, the bank could direct that order first to the exchange that paid the biggest rebate for it.

If the Puzzle Masters were right, and the design of their new exchange eliminated the advantage of speed, it would reduce the informational value of investors’ stock-market orders to zero. If those orders couldn’t be exploited on this new exchange — if the information they contained about investors’ trading intentions was worthless — who would pay for the right to execute them? The big Wall Street banks and online brokers that routed investors’ stock market orders to the new exchange would surrender billions of dollars in revenues in the process. And that, as everyone involved understood, wouldn’t happen without a fight.

Their new exchange needed a name. They called it the Investors Exchange, which wound up being shortened to IEX. Before it opened, on Oct. 25, 2013, the 32 employees of IEX made private guesses as to how many shares they would trade the first day and the first week. The median of the estimates came in at 159,500 shares the first day and 2.75 million shares the first week. The lowest estimate came from the only one among them who had ever built a new stock market from scratch: 2,500 shares the first day and 100,000 the first week. Of the nearly 100 banks and stockbrokerage firms in various stages of agreeing to connect to IEX, most of them

small outfits, only about 15 were ready on the first day. Katsuyama guessed, or perhaps hoped, that the exchange would trade between 40 and 50 million shares a day by the end of the first year — that's about what IEX needed to trade to cover its running costs. If it failed to do that, there was a question of how long it could last. Katsuyama thought that their bid to create an example of a fair financial market — and maybe change Wall Street's culture — could take more than a year. And, he said, "It's over when we run out of money."

On the first day, IEX traded 568,524 shares. Most of the volume came from regional brokerage firms and Wall Street brokers that had no dark pools — RBC and Sanford C. Bernstein. The first week, IEX traded a bit over 12 million shares. Each week after that, the volume grew slightly, until, in the third week of December, IEX was trading roughly 50 million shares each week. On Wednesday, Dec. 18, it traded 11,827,232 shares. But IEX still wasn't attracting many orders from the big banks. Goldman Sachs, for example, had connected to the exchange, but its orders were arriving in tiny lot sizes, resting for just a few seconds, then leaving.

The first different-looking stock-market order sent by Goldman to IEX landed on Dec. 19, 2013, at 3:09.42 p.m. 662 milliseconds 361 microseconds 406 nanoseconds. Anyone who was in IEX's one-room office when it arrived would have known that something unusual was happening. The computer screens jitterbugged as the information flowed into the market in an entirely new way — lingering there long enough to trade. One by one, the employees arose from their chairs. Then they began to shout.

"We're at 15 million!" someone yelled, 10 minutes into the surge. In the previous 331 minutes they had traded roughly 14 million shares.

"Twenty million!"

"Thirty million!"

"We just passed AMEX," shouted John Schwall, their chief financial officer, referring to the American Stock Exchange. "We're ahead of AMEX in market share."

"And we gave them a 120-year head start," Ryan said, playing a little loose with history. Someone had given him a \$300 bottle of Champagne. He'd told Schwall that it was worth only \$100, because Schwall didn't want anyone inside IEX accepting gifts worth more than that from outsiders. Now Ryan fished the contraband from under his desk and found some paper cups.

Someone put down a phone and said, "That was J. P. Morgan, asking, 'What just happened?' They say they may have to do something."

Someone else put down a phone. "That was Goldman. They say they aren't even big. They're coming big tomorrow."

J. P. Morgan, in other words, might actually route investors' trades to IEX, and Goldman might route more of them than they had done so far.

"Forty million!"

Fifty-one minutes after Goldman Sachs gave them their first honest shot at Wall Street customers' stock-market orders, the U.S. stock market closed. Katsuyama walked off the floor and into a small office, enclosed by glass. He thought through what had just happened. "We needed one person to buy in and say, 'You're right,'" he said. "It means that Goldman Sachs

agrees with us. Now the others can't ignore this. They can't marginalize it." Then he blinked and said, "I could [expletive] cry now."

He'd just been given a glimpse of the future — he felt certain of it. If Goldman Sachs was willing to acknowledge to investors that this new market was the best chance for fairness and stability, the other banks would be pressured to follow. The more orders that flowed onto IEX, the better the experience for investors and the harder it would be for the banks to evade this new, fairer market.

IEX had made its point: That to function properly, a financial market didn't need to be rigged in someone's favor. It didn't need payment for order flow and co-location and all sorts of unfair advantages possessed by a small handful of traders. All it needed was for investors to take responsibility for understanding it, and then to seize its controls. "The backbone of the market," Katsuyama says, "is investors coming together to trade."

Three weeks later, two months after IEX opened for business, 14 men — the chief executives or the head traders of some of the world's biggest money managers — gathered in a conference room on top of a Manhattan skyscraper. Together they controlled roughly \$2.6 trillion in stock market investments, or about 20 percent of the U.S. market. They had flown in from around the country to hear Katsuyama describe what he learned about the U.S. stock market since IEX opened for trading. "This is the perfect seat to figure all this out," Katsuyama said. "It's not like you can stand outside and watch. We had to be in the game to see it."

What he had discovered was just how badly the market wanted to remain in the shadows. Despite Goldman's activity, many of the big banks were not following the instructions from investors, their customers, to send orders to IEX. A few of the investors in the room knew this; the rest now learned as much. One of them said: "When we told them we wanted to route to IEX, they said: 'Why would you want that? We can't do that!'" After the first six weeks of IEX's life, a big Wall Street bank inadvertently disclosed to one big investor that it hadn't routed a single order to IEX — despite explicit directions from the investor to do so. Another big mutual fund manager estimated that when he told the big banks to route to IEX, they had done so "at most 10 percent of the time." A fourth investor was told by three different banks that they didn't want to connect to IEX because they didn't want to pay their vendors the \$300-a-month connection fee.

Other banks were mostly passive-aggressive, but there were occasions when they became simply aggressive. Katsuyama heard that Credit Suisse employees had spread rumors that IEX wasn't actually independent but owned by RBC — and thus just a tool of a big bank. He also heard what the big Wall Street banks were already saying to investors to dissuade them from sending orders to IEX: The 350-microsecond delay IEX had introduced to foil the stock-market predator made IEX too slow.

Soon after it opened for trading, IEX published statistics to describe, in a general way, what was happening in its market. Despite the best efforts of Wall Street banks, the average size of IEX's trades was by far the biggest of any stock exchange, public or private. Trades on IEX were also four times as likely as those elsewhere to trade at the midpoint between the current market bid and offer — which is to say, the price that most would agree was fair. Despite the reluctance of the big Wall Street banks to send orders to IEX, the new exchange was already making the dark pools and public exchanges look bad, even by their own screwed-up standards.

Katsuyama opened the floor for questions.

“Do you think of [high-frequency traders] differently than you did before you opened?” someone asked.

“I hate them a lot less than before we started,” Katsuyama said. “This is not their fault. I think most of them have just rationalized that the market is creating the inefficiencies, and they are just capitalizing on them. Really it’s brilliant what they have done within the bounds of the regulation. They are much less of a villain than I thought. The system has let down the investor.”

“How many good brokers are there?” asked an investor.

“Ten,” Katsuyama said. (IEX had dealings with 94.) The 10 included RBC, Bernstein and a bunch of even smaller outfits that seemed to be acting in the best interests of their investors.

“Three are meaningful,” he added: Morgan Stanley, J. P. Morgan and Goldman Sachs.

One investor asked, “Why would any broker behave well?”

“The long-term benefit is that when the [expletive] hits the fan, it will quickly become clear who made good decisions and who made bad decisions,” Katsuyama said. In other words, when some future stock-market crash happened, perhaps a result of the market’s technological complexity, the big Wall Street banks would get the blame.

The stock market really was rigged. Katsuyama often wondered how enterprising politicians and plaintiffs’ lawyers and state attorneys general would respond to that realization. (This March, the New York attorney general, Eric Schneiderman, announced a new investigation of the stock exchanges and the dark pools, and their relationships with high-frequency traders. Not long after, the president of Goldman Sachs, Gary Cohn, published an op-ed in *The Wall Street Journal*, saying that Goldman wanted nothing to do with the bad things happening in the stock market.) The thought of going after those who profited didn’t give Katsuyama all that much pleasure. He just wanted to fix the problem. At some level, he still didn’t understand why some Wall Street banks needed to make his task so difficult.

Technology had collided with Wall Street in a peculiar way. It had been used to increase efficiency. But it had also been used to introduce a peculiar sort of market inefficiency. Taking advantage of loopholes in some well-meaning regulation introduced in the mid-2000s, some large amount of what Wall Street had been doing with technology was simply so someone inside the financial markets would know something that the outside world did not. The same system that once gave us subprime-mortgage collateralized debt obligations no investor could possibly truly understand now gave us stock-market trades involving fractions of a penny that occurred at unsafe speeds using order types that no investor could possibly truly understand. That is why Brad Katsuyama’s desire to explain things so that others would understand was so seditious. He attacked the newly automated financial system at its core, where the money was made from its incomprehensibility.

----- ///

This article is adapted from “Flash Boys: A Wall Street Revolt,” by Michael Lewis, published by W. W. Norton & Company. Michael Lewis is the author of “Boomerang” and “The Big Short.”

Editor: Dean Robinson ... A version of this article appears in print on April 6, 2014, on page MM26 of the *Sunday Magazine* with the headline: *The Wolf Hunters of Wall Street*.

----- ///

Contact information for Peter Burgess: Founder / CEO ... TrueValueMetrics

Website: <http://truevaluemetrics.org>

Email: peterbnyc@gmail.com

Skype: peterburgessnyc Twitter: @truevaluemetric

LinkedIn for Peter Burgess : www.linkedin.com/in/peterburgess1/

This paper as PDF:

<http://www.truevaluemetrics.org/DBpdfs/MDIA/Ideas-Stock-Market-Dysfunction-Exposed.pdf>