Goldman
Sachs

The 'Outsiders' - Emerging Ecosystems

Platforms of Change

From Grassroots to Moonshots

Year one of a new administration and the headlines have kept us busy. Market rotations have been swift, the rise of passive unabated and the distinction between quant and fundamental continues to fade. While this construct remains our reality, we pause for a moment to look beyond the benchmark and peer into seven emerging platforms that exist in the periphery of our investable universe yet merit close watching.

From **DIY algorithms** broadening the pool of alpha to the multi-billion dollar **Dark Net** economy and the rise of **Voice & Digital Assistants in Search**, industries we know are silently reshaping. The blending of science, technology and **Sports Analytics** is changing the nature of engagement between coach, athlete and fan, while the dual forces of AI and the IoT are shifting the IT pendulum back away from the cloud towards **Edge Computing**. Nature's deeper questions, including the search for a cure for **Alzheimer's** and finding order in disorder via **Chaos Theory**, may have once seemed out of reach but are increasingly within our grasp.

Rise of the New

(1) DIY Algorithms: Democratizing Finance. Dorm rooms to desktops the crowdsourcing of alpha may be the next investment paradigm.

(2) Edge Computing. The IT pendulum is once again shifting from centralized (i.e., cloud-based) to decentralized computing in end-devices.

(3) Big Data & Sports Analytics. From optical tracking to next-gen wearables, the marriage of science and sport is increasingly ubiquitous.

(4) Is Voice the Death of Search? From Siri to Cortana, voice search could shake up the $100bn+ advertising market while also driving its growth.

(5) The Hidden Order of Chaos Theory. What if complete disorder is mathematically impossible?

(6) Alzheimer's in Focus. An est. 5.2mn people in the US have AD. With 350 active clinical trials and strong VC investment, hope for a cure abounds.

(7) The Dark Net 101. Anonymous browsing using tools created by the US government fosters a multibillion dollar shadow economy.

Robert D. Boroujerdi
(212) 902-9158 robert.boroujerdi@gs.com
Goldman Sachs & Co. LLC

Christopher Wolf, CFA
(212) 934-4221 christopher.wolf@gs.com
Goldman Sachs & Co. LLC

Table of Contents

Analyst Team Contributors

Robert D. Boroujerdi
'DIY Algorithms: Democratizing Finance'
robert.boroujerdi@gs.com
212-902-9158
Goldman Sachs & Co. LLC

Heather Bellini, CFA
'Is Voice the Death of Search?'
heather.bellini@gs.com
212-357-7710
Goldman Sachs & Co. LLC

Terence Flynn, PhD
'Alzheimer's In Focus'
terence.flynn@gs.com
212-357-5057
Goldman Sachs & Co. LLC

Simona Jankowski, CFA
'Edge Computing' & 'Is Voice the Death of Search?'
simona.jankowski@gs.com
415-249-7437
Goldman Sachs & Co. LLC

Heath P. Terry, CFA
'Is Voice the Death of Search?' & 'The Dark Net 101'
heath.terry@gs.com
212-357-1849
Goldman Sachs & Co. LLC

Christopher Wolf, CFA
'Big Data and Sports Analytics: Not Just Moneyball'
christopher.wolf@gs.com
212-934-4221
Goldman Sachs & Co. LLC

Dan Duggan
'The Hidden Order of Chaos Theory'
dan.duggan@gs.com
212-902-4726
Goldman Sachs & Co. LLC

1) DIY Algorithms: Democratizing Finance

Robert D. Boroujerdi

The democratization of finance is happening all around us. Anything analog is being challenged by the digital as the twin forces of lower compute power and increased machine learning capabilities further pervade capital markets. We have long written about the 'Rise of the Machines' in the form of ETF, factor investing and quantitative strategies, but an emerging business model stands to potentially disrupt the disrupters. Indeed, while change on Wall Street is constant the emergence of DIY Algorithmic Investing aims to level the playing field if not challenge traditional convention in the world of quantitative investing. These crowdsourced alternative hedge fund models are seeing their customers shift from hobbyists on their couches slinging code to a host of global professionals in fields ranging from data science to programming to engineering leveraging curiosity and analytics to create profit. Aided by the carrot of payouts on performance that exceed traditional money managers to the potential for institutional backing and in many cases no loss of their IP to a third party, this emerging ecosystem bears watching as asset management becomes more technology-centric.

The Democratization Begins. Leveraging a web-based platform, DIY algorithmic trading companies allow an individual with coding expertise (e.g., Python) the ability to 1.) create, 2.) back-test, and 3.) run rule-based strategies against actual live stock market activity. At their disposal stands an ever-expanding suite of historical data ranging from financial ratios to historical prices to technical signals. Leveraging this data, individuals can curate, program and profit from algorithms aimed at everything from momentum to mean-reversion to trendline strategies. Most importantly, the individual is empowered without heavy overhead both in the forms of hours of work and upfront costs. Quantopian, a leader in the space, has seen material user growth over the last 5 years while financial python packages are heavily skewed to analyzing market data and creating algorithms (Exhibits 1 and 2).

Exhibit 1: Quantopian's user base has grown to more than 140,000 members
Quantopian members; 2012 – 2017 Q2

**Exhibit 2: Financial Python packages are skewed towards Market Data & Algos**
Percentage of Available Python Financial Packages (by Category)

Source: Quantopian, Goldman Sachs Global Investment Research.

Source: Python.org, Goldman Sachs Global Investment Research.

Exhibit 3: How Stuff Works 101: A DIY Algorithm Flow Chart
Illustrative Example of The Birth and Prosecution of an Algorithm
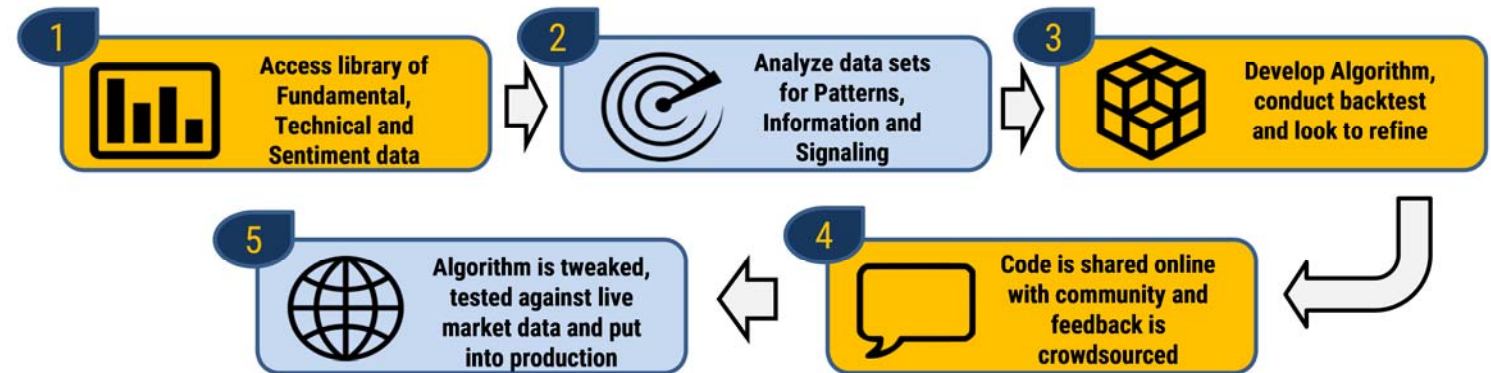
Source: Goldman Sachs Global Investment Research.

The power of the Crowd. Join any discount or full scale broker and a host of data sits at your fingerprints. From quotes to news to research to charting tools, the professional investor toolkit has grown exponentially over the last few decades. But outside a discussion with your broker, a visit to the FAQ page or a chatroom much of an individual investor's experience happens in isolation. However, much like open-source software, the creator of a DIY algorithmic strategy can allow other individuals on the platform to edit, enhance or alter the original code depending on the platform. This community angle opens up vast sharing networks where open-sourced libraries of bits of code are made available to those just starting or professionals who are tweaking existing strategies. The sites often also offer up tutorials for those new to investing or even coding. More powerful, perhaps, is the significant global reach of those involved.

Institutionalizing the Individual. One of the key areas of DIY algorithms that have gained attention is the potential for institutional money to back algorithmic trading strategies. Be it from trading contests (e.g., data-thons) where the "best" algorithm is funded, to hedge funds holding competitions (e.g., Kaggle) to find and hire talent the race to find alpha is going mainstream. In this way the group certainly can take the form of a new type of asset manager where parameters of strategy are clearly delineated (e.g., strategies with defined turnover, market beta, volatility, etc.). Clearly, from a recruiting standpoint never have Wall Street and Silicon Valley been more locked in a war of talent where this generation of recent graduates are focused on mobility, experiences and freedom from traditional convention.

Considerations. From the institutional investor backing case the pros and cons are worth watching. For instance, on the positive side capacity in these strategies potentially could be more open while also improving transparency. With that said many pensions, endowments and foundations may want a more traditional, conservative approach. Further, like most things algorithms look best on paper and may not translate to the real world (e.g., overfitting a back-test, capacity constraints in real world trading, survivorship bias in studies). Thus the journey from inception to action can take a long time to incubate with high failure rates. In addition, like any strategy, the more people do it the less efficacy in the strategy remains … not to mention the potential for how an algorithm acts in a tail-risk event.

Sample companies to watch

Quantopian. Boston-based online platform that enables freelance developers to build and test algorithmic trading strategies and realize a share of the profits when the strategy is put into production and works. Over the last four years the Quantopian community has grown to 140,000 members and created 500,000+ algorithms. Developers benefit from not having to raise capital or manage their code, all while retaining their intellectual property on a crowd-sourced platform with online educational modules. In April 2017, Quantopian started managing investor capital with allocations ranging from $100,000 to $3,000,000 per algorithm. The company aims, by the end of 2017, to make allocations that average $5-$10mn per algorithm with a minimum of $1mn and a high of $50mn.

Numerai. This San Francisco-based hedge fund converts financial data into abstract form where anonymous data scientists use AI to create machine learning models. Developers are incentivized through competitions paid in Numeraire, Numerai's own cryptographic token on the Ethereum blockchain. Each competition allows Numerai to add new machine learning models into one expanding meta model. With the meta model investors do not participate in the upside, but with a hard dollar check for the most useful/impactful submission. In many ways this platform differs in that it offers more of a challenge to sift through a data puzzle.

Quantiacs. Web-based "marketplace" for quants to develop algorithms, publish performance and connect with investors. Quantiacs provides the platform and data required for users to generate and test their algorithms and splits the 20% performance fees they charge investors 50/50 with the algo's creator. Developers can enter Quantiacs-funded competitions where the three best trading strategies are allocated $1,000,000, $750,000 and $500,000, respectively at this time. While the company invests with its own money at this time they also focus on providing a marketplace for others to find talent and access.

Cloud9Trader. Cloud-based platform still in beta pre-release that enables quants to develop and backtest trading strategies with computationally intensive tick-level precision. Offers a paid subscription model in exchange for full database access coupled with the ability to send trade execution signals to your broker.

QuantConnect. Cloud-based platform with over 34,000 users where quants can develop, back test, and deploy encrypted strategies live to their brokerage accounts. Subscription-based model offers users access to 400TB tick-level data library, ability to code in multiple languages and co-located market servers for live algorithms.

2) Edge Computing: Move over Cloud Computing… the IT pendulum is starting to swing

Simona Jankowski, CFA

Over the past decade, cloud computing has become the dominant paradigm in IT, giving rise to a new set of winners (e.g. Amazon) and losers (legacy IT vendors). Just as complacency has set in and portfolios have been aligned to this new world order, the IT pendulum is set to swing again toward the emerging paradigm of "edge" (or "fog") computing, fueled by the twin forces of the Internet of Things (IoT) and Artificial Intelligence (AI). With the advent of IoT, the amount of data generated at the "edge" (i.e., in the devices themselves – e.g., autonomous cars, robots, smart cities) will simply be too large, and in some cases require too rapid of a response (e.g., tap the brakes) to be sent back for processing in the cloud. Instead, the processing of that data will happen locally in the device and in nearby servers/infrastructure (e.g., on the factory floor).

Over the decades, Information Technology (IT) has behaved somewhat like a pendulum – swinging between two extremes: centralized and decentralized architecture. A centralized architecture (think cloud computing) does most of the data crunching function in a central location (e.g., Amazon's data center), which can be accessed by remote devices (e.g., phones, PCs). On the other hand, in a decentralized architecture, most of the data crunching happens in the devices themselves (think back to the big PC tower under your desk a decade ago). The early days of IT (1950s through the 1970s) were defined by mainframes, which were centralized. Then the pendulum swung to decentralized architecture with the client-server era of the 1980s-early 2000s, which put a desktop on every desk and a "computer closet" on every office floor. With the advent of fast Internet connections, the pendulum swung back again to centralized IT in the form of cloud computing, which has been the dominant IT trend of the last decade.

Exhibit 4: The IT pendulum is about to swing again from centralized to decentralized computing
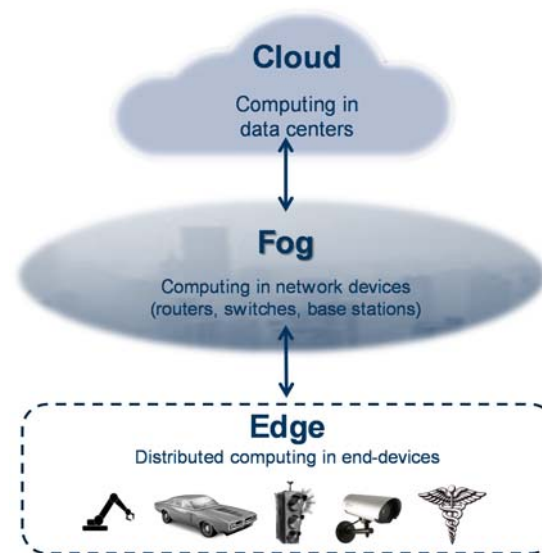Evolution of enterprise computing
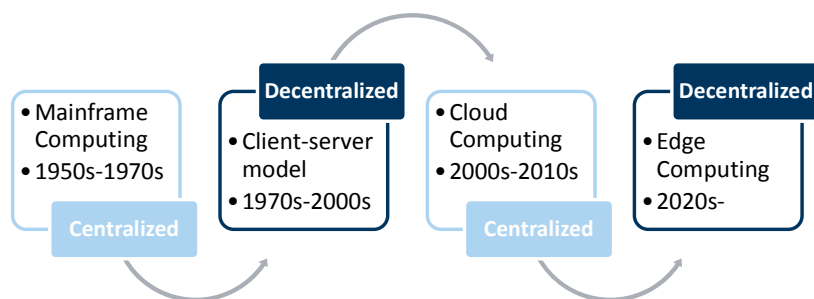
Exhibit 5: The emerging architecture for IoT and AI
Cloud, Fog and Edge computing working together

Source: Goldman Sachs Global Investment Research.

Source: Goldman Sachs Global Investment Research.

Living on the Edge: why now? So why is IT moving to Edge Computing when it feels like Cloud Computing is still in the early-to-mid innings of adoption? First off, it's important to note that Edge Computing will supplement, not replace, Cloud Computing. The twin forces driving the adoption of Edge Computing are the Internet of Things and Artificial Intelligence.

1) **The Internet of Things** – IoT requires distributed intelligence and analytics in addition to hosting those capabilities in the cloud. This is because much of the data generated by IoT devices merely needs to be monitored for events, rather than processed in its entirety. For example, a video camera capturing footage in a city parking lot only needs to alert the police or fire department when there is an event, such as suspicious activity or fire. It is not cost efficient (from a bandwidth usage perspective) to carry all that video traffic 3,000 miles away to a cloud data center if the processing of the data can be done locally. There could be an additional latency reason for processing the data in the edge rather than the cloud – especially when the response to an event has to be instantaneous, such as slamming on the brakes of a train in response to an obstacle on the tracks.

2) **Artificial Intelligence** – The AI renaissance of the last few years has been fueled by new machine learning techniques, which require teaching computers what to do (training), and then having the computers apply that training to complete certain tasks (inference). Training is very compute-intensive and is typically done in a data center. However, inference is often applied in the "edge" of the network, when a decision needs to be made based on data collected at the edge. For example, in the above example of the city parking lot, the video camera footage can be augmented by sensors that capture inputs such as noise, pressure, temperature, etc., which can be analyzed with the use of AI models to infer whether there is a gunshot, fire, or other emergency.

Fog vs. Edge Computing – what is the difference? The terms "Edge Computing" and "Fog Computing" are sometimes used interchangeably, but there is a distinction. Cisco coined the term "fog" computing, a complement to cloud computing, to refer to the distributed intelligence and analytics in infrastructure at the edge of the network, such as routers, switches, gateways and servers. By contrast, edge computing refers to intelligence and processing power in the end devices themselves (e.g., autonomous cars, robots, etc.), though in this report we have used it more broadly to designate the overall trend.

Who benefits from edge computing? Traditional hardware vendors (e.g., **CSCO, HPE, Dell/EMC, NTAP**) have been most hurt by cloud computing adoption as the major cloud providers tend not to use branded IT hardware. Conversely, we expect them to have more opportunities as the pendulum swings to edge computing and enterprises will need to adopt more computing, storage and networking hardware on their premises.

Within the context of the Semiconductor space, we would expect key beneficiaries to be: **NVDA** (with their Tegra processors and Jetson embedded products for inference), as well as **INTC** (through its Movidius and Nervana products) and **AMD**, to a lesser extent. Analog and sensor companies including **TXN, ADI, MXIM,** and **IDTI** will also benefit, in our view, as they play a key role in converting the real (analog) world into digital data that can be analyzed for insights.

Amazon Web Services (AWS) is also uniquely positioned to deliver compute at the edge given its pace of innovation and market share in centralized cloud computing services. In particular, AWS Greengrass enables local compute, messaging, and synching capabilities for IoT devices. Leveraging its serverless compute service AWS Lambda, AWS Greengrass allows devices to quickly respond locally, even with intermittent connection, and extends the benefits of AWS (management, analytics, storage) to connected devices. The ability to message between devices without connection to AWS lowers costs as well, by filtering locally and only transmitting necessary data back to the cloud. The company is also growing its set of AI-enabled devices at the edge of the network with increasingly sophisticated processing capabilities through the proliferation of Amazon Echo, Echo Dot, Echo Look, and Echo Show. Simply put, Amazon's abilities in cloud computing position it well to see edge computing augment its core AWS offerings.

MSFT CEO Satya Nadella has also noted that his company will be one of the key beneficiaries of edge computing. Microsoft's strengths in both what it terms the "Intelligent Cloud" and the "Intelligent Edge" uniquely position the company to help companies pair cloud computing and edge computing. For instance, with the announcement of Azure IoT Edge, cloud functionality can be exported, packaged, and run on IoT devices remotely without the involvement of a cloud loop, enabling customers to make faster decisions while reducing bandwidth costs. Microsoft believes that big data has gravity, and as a result, computational power will move toward where data is located. Furthermore, as more customers leverage AI (Microsoft now offers 29 distinct cognitive services), edge computing gives customers the ability to train models in the cloud (taking advantage of on-demand compute) and run models in the edge, where the data is physically located. Microsoft's strength in AI also allows the Intelligent Edge to become the interface to the real world. For instance, Microsoft's object recognition cognitive services allow searching for a physical object via a search engine. Overall, we believe that Microsoft's unique positioning in both on-premise and in public cloud will drive incremental demand for its offerings as edge computing workloads grow.

Providers of telecom infrastructure benefit from edge computing owing to the need for high-capacity, low-latency network connections close to end users. This includes a need for more dark fiber provided by companies such as **ZAYO**, wireless small cell systems operated by companies such as **CCI**, and highly interconnected data centers operated by companies such as **EQIX**.

3) Big Data and Sports Analytics: Not Just Moneyball

Christopher Wolf, CFA

The prevalence of big data, wearables, player tracking and advanced analytics in sports has passed the tipping point as adoption rates rise and the breadth of use cases grows. Indeed, the insatiable thirst for a competitive edge coupled with new technologies and advanced computing power (e.g., IBM Watson) is driving the evolution of sports from the analog to the digital era. While the narrative of Bill James, Billy Beane, sabermetrics and "Moneyball" may have shed light on the merit of data over intuition in athletics, Major League Baseball is not where the story ends. Coaches, teams and athletes ranging from Manchester United to your local varsity basketball squad are increasingly incorporating player tracking, wearables, video and statistical analysis into their daily routines. While hype is high, demand growing and the potential addressable market large, competition is already rife across the value chain. That said, the natural marriage of science and sport is only strengthening and in the following pages we address what's driving this renaissance and why it matters.

What is "Sports Analytics"? Perhaps best described as the intersection of science, technology and sport, the ability to capture, track and analyze new categories of data is ushering modern athletics out of the dark ages. At its core, sports analytics is the use of this data to assess, inform and improve team or player performance. While not all facets of sports analytics are dependent on new technologies, teams today are embracing the use of GPS trackers, accelerometers, biometric sensors and advanced optical player tracking to create a wealth of new information that can be reviewed, scrutinized and analyzed in search of a competitive edge.

Why does it matter? Where coaches and athletes once relied on experience, perception and instinct as the primary inputs to their decision making, they now also have empirical data and analysis at their disposal. But transforming a data dump into a treasure trove of insights requires time, money and resources (equipment and personnel), raising the question, can your team keep up? Meanwhile, an entire ecosystem is flourishing around the creation, tracking, distribution and analysis of data that extends far beyond the competitors themselves. From fantasy leagues, sports gambling platforms, full-time analysts, video games, casual fans and even academics, the scope of data consumption is broad and demand accelerating (as are the associated profit pools). Further, we see tailwinds from the consumer's adoption of 'wearables' and 'lifestyles of health and sustainability' (LOHAS) as products initially targeting professional athletes trickle down-market to the far deeper pool of amateur-level competitors. Lastly, not to be overlooked are the potential benefits to player health and safety, a rising concern across all levels of sport. Athlete tracking and biometric data can help identify signs of excessive exertion, stress or fatigue and play a crucial role in injury prevention or recovery.

Exhibit 6: How big is the business of Professional Sport?
Aggregate Franchise value and Annual Revenues of Select Teams

	Football	Baseball	Basketball	Soccer*	Hockey	
	NFL	MLB	NBA	Various	NHL	**Total**
Aggregate Franchise Value	$75.3b	$38.8b	$37.4b	$29.6b	$15.6b	**$196.7b**
Aggregate Revenue	$12.2b	$8.4b	$5.2b	$8.2b	$4.1b	**$38.0b**
# Teams	32	30	30	20	30	**142**
# Players (est.)	1,696	862	449	500	707	**4,214**

*Note: Soccer represents a composite of 20 of the top European clubs.

Source: Forbes Sports Money Index, Goldman Sachs Global Investment Research.

Exhibit 7: The N. America Professional Sports Market: $64bn and Growing
North America Sports Market by Segment

Source: PwC Sports Outlook (October 2016)

Moving beyond "Moneyball": A case study in advanced basketball analytics. If baseball pioneered advanced analytics, the NBA has been an early adopter of the digital era. Ahead of the 2013-2014 season, in every arena that didn't already have them, the league installed STATS SportVu® cameras, an advanced optical tracking system that captures ball and player movements 25 times per second. This data set, used by every NBA team, was made available to a team of researchers from Harvard University who in turn demonstrated its potential power in a paper presented at the 2014 MIT Sloan Sports Analytics Conference.[1] Their computationally intensive, probabilistic framework combines inputs such as the tendencies and skill of the player handling the ball and the positioning of other players on the court to determine an "Expected Possession Value (EPV)". This metric quantifies the number of points the offensive team is expected to score throughout each moment of their possession based on the weighted average probability of all future potential outcomes. Such a fluid framework better mirrors basketball's free-flowing play and opens the door for a host of derived metrics such as "EPV-Added" (i.e., does Player X make better decisions than the average player?). Below we walk through a hypothetical example of how it works (Exhibit 8).

Exhibit 8: New metrics like "Expected Possession Value" advance our understanding of the dynamic sport of basketball
Conceptual and illustrative example of the change in "Expected Possession Value" through a hypothetical basketball possession
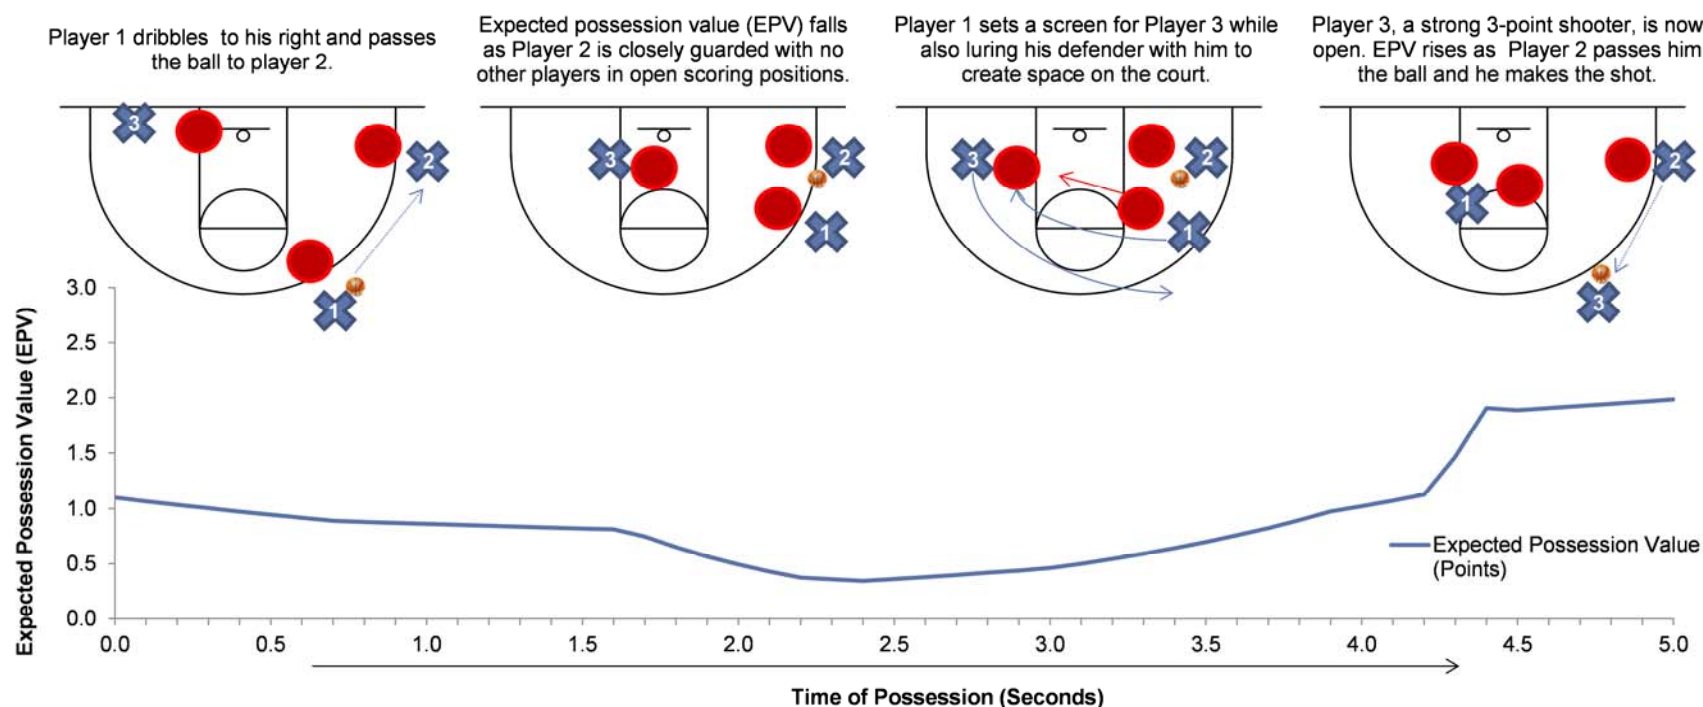
Source: Goldman Sachs Global Investment Research.

---

[1] *POINTWISE: Predicting Points and Valuing Decisions in Real Time with NBA Optical Tracking Data;* Dan Cervone, Alexander D'Amour, Luke Bornn, and Kirk Goldsberry; Department of Statistics and Institute for Quantitative Social Science; Harvard University

Sizing the sports economy... What's at stake? The size of the professional sports industry and financial stakes involved in are significant. Based upon data aggregated from Forbes, the aggregate franchise value of the collective NFL, NHL, NBA and MLB teams combined with 20 of the top European soccer clubs is nearly $200 billion, or roughly equivalent to the GDP of Portugal. The $38 billion in annual revenues these teams collectively generate is in line with the annual sales of Coca-Cola, Honeywell or Oracle (see Exhibit 6). In their latest Sports Outlook report, PricewaterhouseCoopers sizes the North American Sports Market at $64bn in total gate revenues, media rights, sponsorship and merchandising; a figure that is forecasted to grow to $76bn by 2020 (see Exhibit 7).

Further, *Sporting Intelligence*'s annual "Global Sports Salaries Survey" identifies $18.625 billion in aggregate wages for 9,776 athletes participating in 17 different professional sports leagues across 13 countries. This translates to an average of approximately $1.9 million per athlete per year. We acknowledge that estimates of the size of the sports economy can vary considerably, and that the share of spend on sports analytics is just one small part of the broader ecosystem. Still, these figures highlight why it's difficult to find a professional sports franchise that hasn't invested to at least some degree in new analytical hardware, software and/or staff. The value of any incremental edge at this level of sport is simply too significant to be ignored, as is the cost of being left behind.

Is this just about professional sport? No. While naturally professional teams and leagues were the early adopters of new (i.e., expensive) technologies, many of the products on the market today are targeting the semi-professional, collegiate, high school or even youth league team. The incentive for this is clear. The pool of competitive athletes grows exponentially as you move "down market". As shown in Exhibit 9, according to the BLS there are just 13,700 professional athletes in the United States while the latest NFHS High School Athletics Participation Survey identified nearly 7.9 million** high school athletes. Clearly, budgets also shrink exponentially, but not all products are prohibitively expensive. Software platforms that enable basic, real-time do-it-yourself stat tracking on your smartphone or tablet, coupled with video analysis and sharing bear a much smaller price tag.

Further, rising concerns about athlete health and safety across all levels of sport may help spur adoption rates of GPS-tracking/biometric wearables that can provide insights on player exertion and fatigue (e.g., distance traveled, speed, acceleration, heart rate analysis, magnitude and quantity of collisions) and even identify when a player's gait has changed as a result of an injury.

Exhibit 9: Untapped Potential? The pool of "competitive" athletes in the US extends well beyond professionals
Estimated count of professional, collegiate and high school athletes in the U.S
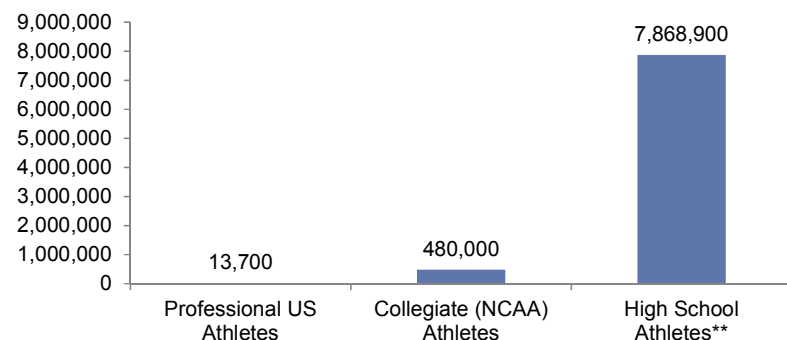
**Note: The NFHS acknowledges that there is some "double counting" of athletes that participate in more than one high school sport.

Source: BLS, NFHS, NCAA, Goldman Sachs Global Investment Research.

Exhibit 10: Rising enthusiasm? Spend on spectator sports is steadily climbing
Consumer Spend on Spectator Sports as % of Total PCE

Spectator Sports as % of Total Personal Consumption Expenditures (PCE)

Source: Bureau of Economic Analysis, Goldman Sachs Global Investment Research.

The ever-growing sports-data ecosystem. Indeed, the use of advanced sports data analytics and statistics is not limited to improving team or player performance. To perhaps state the obvious, one of the primary determinants of the value of a professional sports franchise or league is the size and level of engagement of its fan base. While consumer spending on spectator sports as a percentage of total PCE has been steadily rising for the past 30 years (see Exhibit 10), teams and leagues are constantly competing and seeking new ways to improve fan engagement and in turn drive revenues higher.

As an example of the numerous touch points across the sports-data value chain, the NFL partnered with Zebra Technologies in 2014 to install radio frequency identification (RFID) player tracking in every league stadium (half installed in 2014, half in 2015). This new data set is used by not only the teams themselves, but also the TV broadcasters who can leverage the NFL's "Next Gen Stats" in their commentary (e.g. how much separation was that wide receiver able to create relative to his defender on that last touchdown pass), as well as the fans who can now access and interact with much of this same data on NFL.com. Sportradar, the NFL's consumer data partner, also has the right to distribute this data to its global partners including sport betting platforms which represent a very large and growing source of demand for real-time and advanced sports data and analytics.

Hurdles for adoption and other considerations. As discussed on the prior page, cost and budget limitations coupled with a lack of time or resources to dedicate to data analytics will likely be key headwinds for adoption at the lower levels of competition. While by now many professional teams have dedicated analytics staff and resources, the level of engagement varies by club and by sport. One common refrain from numerous coaches throughout the years at MIT Sloan's Annual Sports Analytics Conference has been a healthy degree of skepticism around accepting data-driven insights at face value. While there is a natural tug-of-war between the value of a coach's life-long experience in a sport versus the bevy of analysts purporting to have new insights and offer new conclusions, this tension may fade as familiarity with new forms of data rises and analytic platforms become more intuitive. Key to this is the ongoing transition from the mere creation of big data sets (e.g., player tracking) to the efficient analysis and use of that data. Lastly, concerns over privacy, ownership and distribution will persist as the breadth of data creation spreads and demand for that data continues to grow.

Meet a few of the innovators and enablers. Already there are dozens of companies dedicated to the creation, analysis and/or distribution of sports data. Below we highlight just a few of the key names and acknowledge that this list is far from exhaustive.

STATS LLC. Established in 1981 and one of the largest sports data and technology providers in the world. Globally they capture data from more than 100,000 events per year. Products offered range from data feeds and live stats/analysis to video analysis platforms, optical player tracking and athlete monitoring. Official statistics provider to Major League Baseball and the PGA, amongst others. STATS SportVu® optical player tracking systems will be used next year by the Ligue de Football Professionnel (LFP) (i.e., top-flight professional soccer in France) and were formerly used by the NBA (through the end of the 2016-17 season).

Sportradar. Official data provider for the NFL, NHL, NBA and NASCAR with global coverage of over 400,000 events annually across more than 40 sports and 800 leagues. Offers data tracking, advanced statistics and live data feeds to clients across broadcasting, media, corporate and sports betting end markets.

Catapult Sports. Offers indoor and outdoor player tracking and monitoring solutions targeted at the elite-level athlete. Products use GPS, GLONASS, or local positioning systems coupled with biometric wearables to measure and track athlete performance. Client base is predominantly elite-level sports programs ranging from Premier League Soccer to NCAA football

Hudl. Provides video review and analysis solutions targeting all levels of competitive sport. Hudl's platform is available for download directly onto a smartphone or tablet and enables coaches to track stats and record/edit video to share or review with their team. It also serves as a platform for communication with team members and a venue for showcasing player talent for recruiters.

4) Voice may put some share "in play", but current leaders likely to re-capture it

Heather Bellini, CFA
Heath P. Terry, CFA
Simona Jankowski, CFA

Tech innovation in the last 5 years has changed voice control from a novelty feature into a legitimate user interface. Apple integrated its digital assistant "Siri" in iOS in 2011. Subsequently, Google, Microsoft, and Amazon have each introduced digital assistants. While early in consumer adoption, we see voice enabling and influencing $100bn in search ad spending, and helping to accelerate the shift of consumer spending online as assistants remove friction from the buying process.

Early in the adoption of voice, commercial use cases are still emerging. For our purposes, "voice" refers to the use of a device (smartphone, smart speaker, etc.) by delivering a vocal command. The result will vary, but could include the delivery of a search engine results page, an audible response to a question, or the operation of other connected devices like home lights. Various surveys have been conducted to determine adoption, and while exact figures vary, the consensus is that adoption skews younger, and the most common use cases are less commercial. Looking up news/weather/traffic, playing music, and smart home controls are universally prevalent uses in survey responses (Accenture, MindMeld, and HigherVisibility).

Voice could impact search ad spending and e-commerce over time. We expect the $100bn global search market to grow at a 13% CAGR through 2020 to reach $146bn. While voice could potentially impact the entire search market over time, we see ads for retail, CPG/consumer products, travel, and consumer electronics (nearly 50% of 2016 search ad spending), as the most likely to be impacted first. As platforms invest to remove friction from the buying process, voice commands are also likely to accelerate the shift in spending to ecommerce vendors, in our view. We note, however, that at this early stage the skills developed for voice-first devices like the Amazon Echo are geared more toward information gathering than commercial activity.

Exhibit 11: The $100bn+ search advertising market could be impacted by adoption of voice inputs/digital assistants over time
US 2016 digital ad spending, by industry

Exhibit 12: Over 65% of the current stable of Alexa skills are focused on fact retrieval or games/trivia/novelty
"Skills" built for Alexa, by category; as of July 2017

Source: eMarketer, Goldman Sachs Global Investment Research

Source: Company data, Goldman Sachs Global Investment Research

Google's speech recognition has improved from a ~75% accuracy rate to 95%+ since 2014

The rapid advancement of voice technology. We view Apple's integration of Siri into iOS in 2011 as the first successful implementation of a voice-enabled digital assistant (VEDA). Google introduced the first iteration of its digital assistant in 2014, followed by Microsoft's Cortana and Amazon's Alexa in 2015. In just the last three years, advancements in neural networks, machine learning, and cloud computing have improved speech recognition and natural language processing capabilities immensely. The speech word error rate for Google's voice technology was about 25% in 2014 and has improved to less than 5% in 2017.

Smartphones were the natural first platform for digital assistants, but we have also seen the evolution of voice-first devices in the home. Since 2014, Amazon has launched a family of voice-first devices under the "Echo" brand, Alphabet has launched its "Google Home" device, and Apple and Microsoft have both announced voice-first devices (the "HomePod" and "Invoke", respectively).

What is the opportunity? We expect the global search market to reach nearly $150bn by 2020, and we expect that the majority of that online search advertising could be served in response to voice queries over time. It is important to note that a voice search will not necessarily generate voice results. Digital assistants on smartphones often display results on screen, accompanied by some audible response. In our research, nearly all of the voice searches we did on smartphones generated on-screen Search Engine Results Pages (SERPs) just as they would have had the query been typed rather than spoken. As more voice-enabled home devices begin to incorporate screens (Amazon's Echo Show/Google Home's integration with Chromecast, for example) we would expect more voice commands to incorporate visual feedback. As such, we may see potential disruption in the search market as default search providers will ultimately vie for control of the ad inventory that appears on screens resulting from voice searches. We see the greatest advantage accruing to the default search providers across the various platforms, as outlined in Exhibit 13.

We believe voice could materially increase overall search volume given the ease with which queries can be made – similar to the expansion seen in PC use following Windows 95 replacing MS-DOS prompts with a more user-friendly robust GUI (graphical user interface). We also see potential transformations of business models for whole industries as some actions like booking online travel may be done primarily via voice in the future versus a keyboard today. We expect most, if not all, mobile applications to incorporate some form of voice-enabled interaction in the future. Just as enterprise software developers today build on a SaaS rather than client-server model, developers are likely to incorporate a "voice-first" or at least a "voice-too" philosophy when developing applications in the future, in our view. Ultimately, a debate related to voice search will be whether or not it will upend the market share of search advertising dollars, which today is dominated by Google with 76% share of the US search ad market in 2016 according to eMarketer, while Bing, Yahoo, and AOL each held single-digit share.

Exhibit 13: Google is the default search provider for popular mobile browsers, but Bing powers 3 of the 4 Digital Assistants

	Digital Assistant				Smartphone Browser		
Platform	Google Home / Assistant	Siri	Amazon Alexa / Echo	Cortana	Chrome	Safari	Samsung Browser
Default Search Provider	Google	Bing	Bing	Bing	Google	Google	Google

Source: Company reports, Goldman Sachs Global Investment Research

Amazon customers spend 10% more on Amazon after purchasing an Echo.

Early read on "voice-first" devices. We see ecommerce as the first commercial beneficiary of voice-enabled smart speakers like the Amazon Echo/Google Home. According to NPD, Amazon Echo owners spent around 10% more on Amazon after purchasing the Echo than they did before. Platform vendors like Amazon and Google are investing to remove friction from the buying process, which we believe will drive more consumers to make regular purchases of less price-sensitive items (paper towels, toothpaste, etc.)

through a voice device over time. The Amazon Echo Show (released on June 28) is the only such device with a screen, and we believe integration with other devices like TVs will play an important role in enabling more commercial activity (booking travel, comparison shopping, etc.) on these devices over time. Google Home has been integrated with Chromecast since its launch, and Amazon introduced a Video Skill API in June to enable Alexa to control smart TVs.

VoiceLabs estimates an installed base of nearly 10mn voice-enabled smart speakers at the end of 2016, and estimates a total installed base of 24.5mn by the end of 2017. These estimates may prove conservative, as they were published prior to Apple's announcement of its HomePod device, which launches later this year. Microsoft will also launch its Cortana-based Invoke device (in partnership with Harmon Kardon) in 2017.

We believe that Voice could have a positive impact on increasing the growth of several categories online in the next 10 years. Based on ecommerce numbers from comScore, consumer packaged goods, a trillion dollar retail category according to the US Department of Commerce, is roughly 2.8% online. CPG represents 8% of the $364 billion US ecommerce market (ex-travel), per comScore. For comparison, Apparel & Accessories, Books & Magazine, and Office Supplies have online penetration of 19.4%, 52%, and 48.4%, respectively.

On Amazon's Echo devices, Alexa defaults to Prime-eligible items from the customer's order history and "Amazon's Choice" items, which Amazon selects in each category based on ratings and price data. For example, when telling Alexa to "buy AA batteries" the "Amazon's Choice" item is its own private label Amazon Basic batteries, which have become one of their highest-selling products in terms of market share. Amazon Basic batteries already account for 1/3 of online battery sales and are seeing 93% YOY growth, and its baby wipes business has gained 16% market share in dollar terms among the top 10 brands, per Nielsen's US Retail Hot Buttons Q1-2017 and techcrunch.com (11/3/2016).

How do you monetize "voice"? Both Amazon and Google have opened up their respective platforms to developers in the hopes of building out a robust ecosystem of third-party voice "apps" (called Skills on Alexa, and Actions on Google Assistant). Neither platform has introduced monetization schema as yet, but we would expect these to eventually follow the same model as we have seen with mobile apps. For example, we expect to see the introduction of paid "skills" on Alexa, or embedded advertising in Google "Actions", within the next year.

What are the risks? The unpredictable nature of consumer behavior remains a risk to voice overall, in our view. We believe it will take time before voice controls make up the majority of search queries, as there are scenarios (offices, classrooms, and various public settings) in which there may be some initial reluctance to use voice commands. However, voice already has high adoption rates among younger users, and we expect overall user acceptance to evolve over time.

5) The Hidden Order of Chaos Theory

Dan Duggan

Tomorrow's winners and losers will be determined not just by how much data they have, but by how much they can do with it. Predictive models are increasingly powered by "exhaust" data – data coming from cell phones, satellite imagery, social media and credit card transactions. As these sources become commoditized the conversation is converging on the algorithms they feed. Terms like "deep learning" and "recurrent neural networks" are now entering the mainstream, reflecting their rising importance. But these are just the latest in a long line of sophisticated algorithms enjoying wide application in diverse areas from portfolio diversification and medicine to electronics design, traffic modeling and equity markets. We take a deeper dive into theories that form the foundation of their impressive power: Chaos and Ramsey theories.

"Chaos" theory is all about order. At its core it describes patterns emerging from apparently random data. We and our traditional algorithms miss these relationships because their behaviors are nonlinear or dynamic. Take the housefly as an example. Its movement may seem chaotic to the untrained eye. But when plotting how its location changes a pattern becomes visible. That path can be represented using Chaos theory. Chaotic models must be tuned to each use case so large amounts of input data are often required. With increasing data sizes comes added risk of misinterpreting those signals (e.g., spurious correlations). This risk increases as datasets become larger, making incorrect assumptions increasingly likely**.** This is where Ramsey theory enters the equation.

Ramsey theory provides the tools to generalize spurious correlations between any number of objects. These tools can be used when drawing conclusions from a study on the length of eyelashes or for portfolio diversification strategy. Its applications are extremely broad with utility to insurance companies, health care providers, portfolio managers or really anyone who cares about risk. **Taken together, Chaos and Ramsey theories provide a road map to solutions of tomorrow's problems.**

Portfolio diversification à la Ramsey theory. Before tackling *Big Data*, let's see how Ramsey theory impacts even *Small Data*. Ramsey theory has a material effect on portfolio optimization for even tiny portfolios. Bear with us as we dig into some technical details in the following example.

- Consider a basic portfolio diversification problem using six stocks (A through F).

- Describe the correlation between stocks as either positive (green) or negative (red).

- Determine whether each of the 20 pairwise correlations is positive or negative.

- Calculate which correlations must exist and are potentially spurious using Ramsey theory.

- Assume the "required" correlations are spurious and assign uncertainty.

After accounting for spurious correlations the variance materially increases (Exhibit 14). The effect grows with the portfolio size, so living with spurious correlations is a fact of life… or in this case, mathematics.


The Flight of the Housefly

**Exhibit 14: The unavoidable effect of spurious correlations**
The spread in expected portfolio returns is underestimated according to Ramsey theory

Source: Goldman Sachs Global Investment Research

**Exhibit 15: Chaos and Ramsey theories in research.**
Their use in academic literature shows their wide range of practical applications

Source: arXiv.org, Goldman Sachs Global Investment Research

In our example a three-stock negative correlation is spurious – why?

- *If A's returns increase*, B's returns must decrease.

- If B's returns decrease, then C's returns must increase.

- If C's returns increases, *then A's returns must decrease.*

- This is a contradiction in a market where these changes are happening simultaneously.

- **At least one of the correlations (A-B, B-C, A-C) must be spurious.**

Ramsey theory reveals that complete disorder is mathematically impossible. In other words, these effects can be partially mitigated but not removed. Suppose we decided that our portfolio can be positively correlated, negatively correlated, or uncorrelated. This would lead to a three-colored system in our first example instead of two. Unavoidable order now emerges in a portfolio with only 17 stocks.

Imagine the hypothetical scenario of a market with only 50 portfolio managers participating. One discovers Ramsey theory. As a result she now has a materially better understanding of her risk. She could adjust her portfolio positions to better match her desired risk profile. Her competitive advantage would remain as long as Ramsey theory stays siloed. Once disseminated, the arbitrage opportunity she currently enjoys and exploits would end.

It's out there…

Studies referencing Chaos and Ramsey theories are on the rise. The recent expansion of research on both indicates the positive potential for future industry use (see Exhibit 15). The 220 approved Chaos-based US patents speak to Chaos theory's longer publication history. Here are some current applications that highlight Chaos theory's diverse utility.

- **Chaos theory in medicine.** Heartbeats contain effects that cannot be modeled by simple periodic behavior. Medical researchers instead use mode-based data representations to understand heart rate variability (HRV). They can extract details related to specific types of cardiac arrest. Other studies investigate models that would use HRV to predict heart failure. Chaos-based patents now even exist for medical implant devices that detect a type of HRV called atrial fibrillation. Combining this technology with health monitoring smartphone apps could make pre-heart attack alerts as common as digital heart rate monitoring.

- **Tackling Moore's law.** Integrated circuits in electronics from laptops to car dashboards continue to rise in popularity. That growing market is now beginning to outpace Moore's Law as transistors reach physical size limitations. Moore's Law says that every two years the number of transistors in an integrated circuit must double to keep pace with demand. To combat these restrictions engineers are instead building chaotic circuits. They do it by leveraging nonlinear chip properties to enable them to multi-task. This drastically drives down computation time and critically reduces noise. These currently exist only in the laboratory but have the potential to disrupt the commercial market.

- **Chaotic traffic patterns.** It may not be surprising that the government was at the cutting edge of Chaos theory over twenty years ago. More surprisingly, the Department of Transportation was one of its pioneers. Their work uses non-linear modeling and different phase space representations to describe traffic flow. These techniques ultimately dictate the timing of stoplights and length of zoned speed limits. More efficient commutes are advantageous for Retail whose revenues are directly tied to footfall. Recurrent neural networks are being tested to create more dynamic descriptions of traffic flow and these have future potential to be finely tuned with the real-time availability of vehicle Big Data.

- **Fractals in action.** Fractals, self-repeating geometric patterns, lie at the very heart of Chaos theory. Though not obvious at first glance, recurrent trends found in fractals are a good fit for the market's cyclic nature. Leading theories like the efficient market theory struggle to explain times of financial stress. In contrast, fractal market theory describes market movements as a mixture of short- and long-term investment timescales. Periods of instability disrupting this balance produce in the model large-scale market shifts like sell-offs or even crises. As the power behind its methods like next-generation algorithms continues to improve its potential to describe and predict these nonlinear trends is on the rise.

The (eventual) democratization of Ramsey and Chaos theory

Will these theories be kept as closely guarded secrets forever? Don't count on it. One barrier preventing their larger adoption is fading away with widespread computing power availability. Open-sourced statistical tools that integrate these techniques like R and Python are increasingly accessible and free to use. Together these tools on a scalable architecture offer a powerful platform to leverage Chaos and Ramsey theories. The challenge becomes instilling intuition within these theories' algorithmic logic. Combined, they could unlock the hidden relationships in data to make the "model of everything" a reality.

6) Alzheimer's In Focus: Aiming to eliminate "long goodbyes"

Terence Flynn, PhD

The arc of biopharma innovation spans preventive vaccines, targeted antibody drugs for a range of diseases, and most recently immuno-oncology (I/O), but has yet to deliver a cure for Alzheimer's disease (AD), viewed as one of the final frontiers of drug development. The likelihood of success (i.e., regulatory approval) for AD drugs has been only 0.4% (from 2002-2012) vs. 10-12% on average for other diseases. The most recent late-stage AD drug failures included LLY's Solanezumab (antibody to amyloid beta) in December 2016 and MRK's Verubecestat (BACE inhibitor) in February 2017. However, companies, investigators and scientists continue to make progress in understanding the underling disease biology (and genetics). There are over 350 active clinical trials of AD drugs, the number of drugs entering early stage development has been relatively consistent (Exhibit 18) and venture capital investment in private AD companies has been strong in 2015 and 2016, although thus far 2017 has been quieter (Exhibit 19). Discovering a true cure for AD is likely to remain elusive over the near term; hence, most efforts are concentrated on identifying disease-modifying treatments, which could open up a $10bn+ market opportunity that could ultimately rival some of the largest therapeutic categories (Exhibit 17).

Why it matters? It is impossible to overstate the growing societal burden of AD, which should make the disease a key public health focus. **Approximately one-third of people aged 85 and older develop AD.** In 2014 an estimated 5.2mn people in the United States had AD, a figure which is expected to increase to nearly 14mn by 2050 (Fargo, Alzheimer's & Dementia 2014; Exhibit 16). **By 2050, one new case of AD is expected to develop every 33 seconds in the United States, or nearly a million new cases per year. The global costs of dementia have increased from $604bn in 2010 to $818bn in 2015 (+35%) and are projected to exceed $1tn in 2018 (World Alzheimer Report 2015).** These costs are driven by direct medical costs, social care costs (professional home/residential/nursing home care) and costs of unpaid care (including by family members; nearly 18bn hours of care in 2013 across 15mn unpaid US caregivers).

Exhibit 16: Aging Baby Boomers will lead to an increase in AD in US


*Source: Adapted from Alzheimer's Association, 2014 Alzheimer's Disease Facts and Figures*

Exhibit 17: WW sales of branded drugs for select diseases

Source: Company data, Goldman Sachs Global Investment Research

Exhibit 18: Phase 1 clinical trials of Alzheimer's drugs

Source: clinicaltrials.gov, Goldman Sachs Global Investment Research

Exhibit 19: VC investment in private AD companies

*Annualized   VC Invested (LH Axis)   # of deals (RH Axis)

Source: CB Insights, Goldman Sachs Global Investment Research

What causes Alzheimer's disease? We don't know for sure, but the brains of people with AD have abnormally high levels of a plaque protein called **amyloid beta (Aβ)**. The "amyloid hypothesis" proposes that accumulation of these amyloid plaques in the brain sets in motion a cascade of toxic events, including the formation of tangles (via a protein called **tau**), breakdown of communication inside and between nerve cells, brain inflammation, and ultimately cell death. This theory is also supported by the occurrence of early-onset Alzheimer's in certain groups genetically predisposed to overproduce Aβ. However, not all researchers agree on the amyloid hypothesis, as there's still debate if the plaques are a driver of disease or just a by-product.

But if the amyloid hypothesis is true, then drugs that can reduce amyloid plaque levels in the brain should be able to slow the progression of AD. Biopharma companies have been focused on developing drugs that bind to and reduce the amyloid plaque (**Aβ antibodies**) or prevent them from forming altogether (**BACE inhibitors**). Another approach that some companies are pursuing involves targeting the **tau** protein, but there's less clinical data on this front. The future of AD treatment may very well involve a combination of multiple drugs with different mechanisms.

Why is AD drug development so challenging? The short answer is that unlike many other diseases, we don't know the root cause of AD. Researchers use animal models to guide drug development before testing in people. However in AD, "efficacy" seen in animals has not been predictive of clinical benefit in people. It's possible that animal models only reflect a single aspect of AD or that they don't accurately represent the progressive nature of the disease, which is characterized by the death of nerves.

Exhibit 20: Select AD drugs in development

Company	Drug	Phase	Target	Next Catalyst
AXON	Intepirdine	3	5-HT6	3Q17: Phase 3 data
Eisai/BIIB	BAN2401	2	Aβ	4Q17: Final Phase 2 data
MRK	Verubecestat	3	BACE	2H17: Analysis of terminated Phase 2/3 study 2018: Phase 3 data from second study
LLY	Lanabecestat	2/3	BACE	2019: Phase 2/3 data
Roche	Crenezumab	3	Aβ	2019+: Interim Phase 3 data
Roche	Gantenerumab	3	Aβ	2019+: Interim Phase 3 data
BIIB/Eisai	Aducanumab	3	Aβ	4Q19/1Q20: Phase 3 data
Eisai/BIIB	Elenbecestat	3	BACE	2020: Phase 3 data
AMGN/NVS	CNP520	2/3	BACE	2023: Phase 2 data

Source: Company data, Clinicaltrials.gov, Goldman Sachs Global Investment Research

What have we learned that could lead to success for the next round of drugs? Past clinical failures suggest that treating moderate stage Alzheimer's disease might be too late in the disease course or that patients without amyloid plaque might be less likely to respond to treatment. Hence we are encouraged that (1) many ongoing AD trials are studying patients with earlier stage disease (prodromal/mild) and amyloid plaque in the brain (measured by PET scan), where it might be easier to demonstrate an effect, (2) some newer drugs lead to greater reductions in amyloid plaque levels vs. older drugs, and (3) combination therapy (antibody to amyloid beta and a BACE inhibitor) might produce better efficacy, although there are no ongoing trials exploring this yet. **There are a number of ongoing mid to late stage clinical trials being conducted by biopharma companies including BIIB (and partner Eisai), Roche and MRK and we expect a steady stream of data (Exhibit 20).**

Clearly a blockbuster opportunity, but how large? There are approved AD drugs that treat the symptoms (not the underlying cause) of the disease but do not prevent disease progression. Despite these limitations, two drugs (Aricept and Namenda) had combined peak sales of nearly $6bn prior to the entry of generic versions. In our view, price and uptake for future disease-modifying drugs could be higher than symptomatic treatments, but this will ultimately be driven by the magnitude of the efficacy benefit.

We previously conducted a sensitivity analysis that suggests a potential US commercial opportunity of $5-8bn for mild AD with Aβ plaque and $3-5bn for prodromal with Aβ plaque (for disease-modifying agents priced at $20k/year and 40-60% share). Hence worldwide sales could exceed $10bn. See *The Alzheimer's conundrum: Putting key issues into perspective ahead of our conference* published April 5, 2016 for more details.

Key risks: Investing in companies developing AD drugs is high risk, as many of the treatments in development are likely to fail as a result of a lack of efficacy and/or acceptable safety/tolerability. Furthermore, even if an AD drug succeeds in clinical trials and reaches the market, the clinical benefit of the drug will be an important variable to consider when evaluating the ultimate market potential.

7) The Dark Net 101

Heath P. Terry, CFA

The "Dark Net" is a term applied to a network of internet sites that allow for anonymous online browsing, communications, and, enabled largely by cryptocurrencies (e.g., Bitcoin), transactions. Though the term lends itself to sensationalism and the actual Dark Net can indeed be used for illegal activities, the original intent behind its creation was simply to enable secure communications with US military bases and other assets in foreign countries.

Today, the Dark Net is still used by political dissidents, whistle blowers, journalism sources wishing to communicate anonymously, and by people living behind national firewalls seeking to reach the outside world. Meanwhile, thousands of sites dedicated to the sale of illicit goods and services operate in a multibillion dollar shadow economy. While this ecosystem sits primarily at the fringe of the internet, the Dark Net's fundamental basis of anonymity and privacy has driven the need for network protocol creativity and other innovations that could carry more mainstream implications over time. Indeed, with a growing set of internet users expressing concerns over privacy, and with tools such as the DuckDuckGo search engine and the Ghostery browser extension rising in awareness, what is born from the Dark Net could ultimately find widespread adoption.

The Dark Net is a small subset of the Deep Web, the part of the internet that isn't indexed by search engines or generally accessible to the public. In contrast to the Surface Web, think of the Deep Web as any site or database that isn't crawled or indexed by a search engine. This includes the directory of the Library of Congress, Census.gov and any site that requires a log in and password, like your email, your bank account, or your company's intranet. The Dark Net goes further in its inaccessibility. It is made up of an estimated 7,000 websites hosted on overlay networks, where the IP address is encrypted into a sixteen character alpha numeric code and then routed through multiple layers of a peer-to-peer network to obscure the location of the site and users.

Exhibit 21: How "Onion Routing" works

In Onion Routing, messages are generated with multiple layers of encryption that are subsequently decrypted or "peeled away" one layer at a time by other TOR nodes until the message arrives unencrypted at its destination.

Source:
TOR client path to
destination server

TOR node
encrypted link
unencrypted link

Source: MakeUseOf.com

Exhibit 22: Who uses the Dark Net? Top 10 countries (based on relay)
Estimated # of directly-connecting clients by country; daily average for trailing 3m

Source: Tor Project

TOR enables the Dark Net. Users primarily access Dark Net sites and communicate with each other using the Tor browser, a free, open-source software that anonymizes the user's location, online history, and communications while allowing access to websites with the .onion domain. The underlying technology used in Tor is called "onion routing," named after the layers of networks that transmissions go through to obscure their origin (see Exhibit 21). Created in the mid-1990s, The Onion Routing Project (TOR) was developed by the US Navy and is currently maintained by The Tor Project, a 501(c)(3) educational nonprofit that is funded by the US State Department, Google, The Electronic Frontier Foundation, and other government and non-government organizations. These organizations support the Dark Net because of the value they see in allowing the free flow of information and communication globally. While the Tor Project and browser are most synonymous with the Dark Net, other platforms such as the Invisible Internet Project (I2P) and Freenet offer similar venues for anonymous browsing.

The Dark Net has a variety of uses. Dark Net abuses that have been prosecuted in western countries include the sale of firearms, illicit substances, stolen data, pirated content, hacking tools and services. Away from the "extreme fringes", the Dark Net is used by journalists and their sources, whistle blowers, intelligence assets, and political dissidents to communicate anonymously. Other users like DoctorX, a physician based in Madrid, provide free, anonymous, substance abuse counselling. Facebook even supports a TOR-enabled version of the site, launched in late 2014. Alternatively, Twitter has limited TOR access as a response to issues with harassment on the service, requiring TOR users to provide a verified phone number to connect to the service.

Silk Road and the Dark Net. The most well-known use of the Dark Net was the development of Silk Road by Ross Ulbricht. Silk Road was an online marketplace for illegal products that started in February 2011. While Ulbricht started the site with financial motivations, he saw Silk Road as a libertarian approach to reducing the dangers of the illegal drug trade and a tool to drive governments to end drug prohibition. By the time it was shut down by the FBI in October 2013 with Ulbricht's arrest and eventual imprisonment, it is estimated that Silk Road facilitated over $1.2bn in sales (largely through Bitcoin) involving 150,000 buyers and 4,000 sellers.

Though the site was shut down with the conviction of Ulbricht, hundreds of smaller versions of the site have come to replace it. Law enforcement agencies continue to target and convict criminals using the Dark Net by developing their own "hacking tools" to try and cut through the layers of anonymity. The FBI alone has successfully prosecuted Dark Net-enabled crimes involving the sale of illegal drugs and firearms and the spread of extremism.

The internet sector could be impacted. While far removed from the publicly traded internet sector, the Dark Net represents a potentially disruptive risk. The presence of high-profile illegal activity could lead to efforts to more directly regulate the internet. Companies that have their own communication services, like Facebook, Twitter, and Google, could easily be swept up in the incremental regulation, limiting their ability to innovate, raising their cost of compliance, and imposing a negative effect on people's use of their services. In parallel, as internet users become more concerned about privacy, they could increasingly use protected services on the Dark Net as a way to circumvent the data collection, targeting, and measurement that many platforms use to maximize the value of their advertising inventory. Indeed, a survey by Pew Research Center found that 86% of adult internet users have taken steps to mask or remove their digital footprint. The Tor Project recently announced a Tor-enabled smartphone prototype designed to circumvent the user data collection done by iOS and Android phones, a critical monetization tool for advertising supported internet companies. From multi-signature escrow payment systems to Bitcoin "tumblers" that further improve the cryptocurrency's anonymity, the innovations occurring primarily at the fringe of the internet could see more mainstream adoption if these privacy concerns intensify.

Reg AC

Each equity analyst's contribution herein was certified under Reg AC by the analyst primarily responsible for that section as follows: I, Name of Analyst, hereby certify that all of the views expressed in this report accurately reflect my personal views about the subject company or companies and its or their securities. I also certify that no part of my compensation was, is, or will be, directly or indirectly, related to the specific recommendations or views expressed in this report.

Each analyst certifies to the views in the sections he/she authored in this report.

Disclosure Appendix

Reg AC

I, Robert D. Boroujerdi, hereby certify that all of the views expressed in this report accurately reflect my personal views about the subject company or companies and its or their securities. I also certify that no part of my compensation was, is or will be, directly or indirectly, related to the specific recommendations or views expressed in this report.

I, Christopher Wolf, CFA, hereby certify that all of the views expressed in this report accurately reflect my personal views, which have not been influenced by considerations of the firm's business or client relationships.

Unless otherwise stated, the individuals listed on the cover page of this report are analysts in Goldman Sachs' Global Investment Research division.

GS Factor Profile

The Goldman Sachs Factor Profile provides investment context for a stock by comparing key attributes to the market (i.e. our coverage universe) and its sector peers. The four key attributes depicted are: Growth, Financial Returns, Multiple (e.g. valuation) and Integrated (a composite of Growth, Financial Returns and Multiple). Growth, Financial Returns and Multiple are calculated by using normalized ranks for specific metrics for each stock. The normalized ranks for the metrics are then averaged and converted into percentiles for the relevant attribute. The precise calculation of each metric may vary depending on the fiscal year, industry and region, but the standard approach is as follows:

Growth is based on a stock's forward-looking sales growth, EBITDA growth and EPS growth (for financial stocks, only EPS and sales growth), with a higher percentile indicating a higher growth company. **Financial Returns** is based on a stock's forward-looking ROE, ROCE and CROCI (for financial stocks, only ROE), with a higher percentile indicating a company with higher financial returns. **Multiple** is based on a stock's forward-looking P/E, P/B, price/dividend (P/D), EV/EBITDA, EV/FCF and EV/Debt Adjusted Cash Flow (DACF) (for financial stocks, only P/E, P/B and P/D), with a higher percentile indicating a stock trading at a higher multiple. The **Integrated** percentile is calculated as the average of the Growth percentile, Financial Returns percentile and (100% - Multiple percentile).

Financial Returns and Multiple use the Goldman Sachs analyst forecasts at the fiscal year-end at least three quarters in the future. Growth uses inputs for the fiscal year at least seven quarters in the future compared with the year at least three quarters in the future (on a per-share basis for all metrics).

For a more detailed description of how we calculate the GS Factor Profile, please contact your GS representative.

Quantum

Quantum is Goldman Sachs' proprietary database providing access to detailed financial statement histories, forecasts and ratios. It can be used for in-depth analysis of a single company, or to make comparisons between companies in different sectors and markets.

GS SUSTAIN

GS SUSTAIN is a global investment strategy aimed at long-term, long-only performance with a low turnover of ideas. The GS SUSTAIN focus list includes leaders our analysis shows to be well positioned to deliver long term outperformance through sustained competitive advantage and superior returns on capital relative to their global industry peers. Leaders are identified based on quantifiable analysis of three aspects of corporate performance: cash return on cash invested, industry positioning and management quality (the effectiveness of companies' management of the environmental, social and governance issues facing their industry).

Disclosures

Distribution of ratings/investment banking relationships

Goldman Sachs Investment Research global Equity coverage universe

	Rating Distribution			Investment Banking Relationships		
	Buy	Hold	Sell	Buy	Hold	Sell
Global	33%	53%	14%	63%	57%	50%

As of April 1, 2017, Goldman Sachs Global Investment Research had investment ratings on 2,857 equity securities. Goldman Sachs assigns stocks as Buys and Sells on various regional Investment Lists; stocks not so assigned are deemed Neutral. Such assignments equate to Buy, Hold and Sell for the purposes of the above disclosure required by the FINRA Rules. See 'Ratings, Coverage groups and views and related definitions' below. The Investment Banking Relationships chart reflects the percentage of subject companies within each rating category for whom Goldman Sachs has provided investment banking services within the previous twelve months.

Regulatory disclosures

Disclosures required by United States laws and regulations

See company-specific regulatory disclosures above for any of the following disclosures required as to companies referred to in this report: manager or co-manager in a pending transaction; 1% or other ownership; compensation for certain services; types of client relationships; managed/co-managed public offerings in prior periods; directorships; for equity securities, market making and/or specialist role. Goldman Sachs trades or may trade as a principal in debt securities (or in related derivatives) of issuers discussed in this report.

The following are additional required disclosures: **Ownership and material conflicts of interest:** Goldman Sachs policy prohibits its analysts, professionals reporting to analysts and members of their households from owning securities of any company in the analyst's area of coverage. **Analyst compensation:** Analysts are paid in part based on the profitability of Goldman Sachs, which includes investment banking revenues. **Analyst as officer or director:** Goldman Sachs policy generally prohibits its analysts, persons reporting to analysts or members of their households from serving as an officer, director or advisor of any company in the analyst's area of coverage. **Non-U.S. Analysts:** Non-U.S. analysts may not be associated persons of Goldman, Sachs & Co. and therefore may not be subject to FINRA Rule 2241 or FINRA Rule 2242 restrictions on communications with subject company, public appearances and trading securities held by the analysts.

Distribution of ratings: See the distribution of ratings disclosure above. **Price chart:** See the price chart, with changes of ratings and price targets in prior periods, above, or, if electronic format or if with respect to multiple companies which are the subject of this report, on the Goldman Sachs website at http://www.gs.com/research/hedge.html.

Additional disclosures required under the laws and regulations of jurisdictions other than the United States

The following disclosures are those required by the jurisdiction indicated, except to the extent already made above pursuant to United States laws and regulations. **Australia:** Goldman Sachs Australia Pty Ltd and its affiliates are not authorised deposit-taking institutions (as that term is defined in the Banking Act 1959 (Cth)) in Australia and do not provide banking services, nor carry on a banking business, in Australia. This research, and any access to it, is intended only for "wholesale clients" within the meaning of the Australian Corporations Act, unless otherwise agreed by Goldman Sachs. In producing research reports, members of the Global Investment Research Division of Goldman Sachs Australia may attend site visits and other meetings hosted by the issuers the subject of its research reports. In some instances the costs of such site visits or meetings may be met in part or in whole by the issuers concerned if Goldman Sachs Australia considers it is appropriate and reasonable in the specific circumstances relating to the site visit or meeting. **Brazil:** Disclosure information in relation to CVM Instruction 483 is available at http://www.gs.com/worldwide/brazil/area/gir/index.html. Where applicable, the Brazil-registered analyst primarily responsible for the content of this research report, as defined in Article 16 of CVM Instruction 483, is the first author named at the beginning of this report, unless indicated otherwise at the end of the text. **Canada:** Goldman Sachs Canada Inc. is an affiliate of The Goldman Sachs Group Inc. and therefore is included in the company specific disclosures relating to Goldman Sachs (as defined above). Goldman Sachs Canada Inc. has approved of, and agreed to take responsibility for, this research report in Canada if and to the extent that Goldman Sachs Canada Inc. disseminates this research report to its clients. **Hong Kong:** Further information on the securities of covered companies referred to in this research may be obtained on request from Goldman Sachs (Asia) L.L.C. **India:** Further information on the subject company or companies referred to in this research may be obtained from Goldman Sachs (India) Securities Private Limited, Research Analyst - SEBI Registration Number INH000001493, 951-A, Rational House, Appasaheb Marathe Marg, Prabhadevi, Mumbai 400 025, India, Corporate Identity Number U74140MH2006FTC160634, Phone +91 22 6616 9000, Fax +91 22 6616 9001. Goldman Sachs may beneficially own 1% or more of the securities (as such term is defined in clause 2 (h) the Indian Securities Contracts (Regulation) Act, 1956) of the subject company or companies referred to in this research report. **Japan:** See below. **Korea:** Further information on the subject company or companies referred to in this research may be obtained from Goldman Sachs (Asia) L.L.C., Seoul Branch. **New Zealand:** Goldman Sachs New Zealand Limited and its affiliates are neither "registered banks" nor "deposit takers" (as defined in the Reserve Bank of New Zealand Act 1989) in New Zealand. This research, and any access to it, is intended for "wholesale clients" (as defined in the Financial Advisers Act 2008) unless otherwise agreed by Goldman Sachs. **Russia:** Research reports distributed in the Russian Federation are not advertising as defined in the Russian legislation, but are information and analysis not having product promotion as their main purpose and do not provide appraisal within the meaning of the Russian legislation on appraisal activity. **Singapore:** Further information on the covered companies referred to in this research may be obtained from Goldman Sachs (Singapore) Pte. (Company Number: 198602165W). **Taiwan:** This material is for reference only and must not be reprinted without permission. Investors should carefully consider their own investment risk. Investment results are the responsibility of the individual investor. **United Kingdom:** Persons who would be categorized as retail clients in the United Kingdom, as such term is defined in the rules of the Financial Conduct Authority, should read this research in conjunction with prior Goldman Sachs research on the covered companies referred to herein and should refer to the risk warnings that have been sent to them by Goldman Sachs International. A copy of these risks warnings, and a glossary of certain financial terms used in this report, are available from Goldman Sachs International on request.

European Union: Disclosure information in relation to Article 4 (1) (d) and Article 6 (2) of the European Commission Directive 2003/125/EC is available at http://www.gs.com/disclosures/europeanpolicy.html which states the European Policy for Managing Conflicts of Interest in Connection with Investment Research.

Japan: Goldman Sachs Japan Co., Ltd. is a Financial Instrument Dealer registered with the Kanto Financial Bureau under registration number Kinsho 69, and a member of Japan Securities Dealers Association, Financial Futures Association of Japan and Type II Financial Instruments Firms Association. Sales and purchase of equities are subject to commission pre-determined with clients plus consumption tax. See company-specific disclosures as to any applicable disclosures required by Japanese stock exchanges, the Japanese Securities Dealers Association or the Japanese Securities Finance Company.

Ratings, coverage groups and views and related definitions

Buy (B), Neutral (N), Sell (S) -Analysts recommend stocks as Buys or Sells for inclusion on various regional Investment Lists. Being assigned a Buy or Sell on an Investment List is determined by a stock's total return potential relative to its coverage. Any stock not assigned as a Buy or a Sell on an Investment List with an active rating (i.e., a stock that is not Rating Suspended, Not Rated, Coverage Suspended or Not Covered), is deemed Neutral. Each regional Investment Review Committee manages various regional Investment Lists to a global guideline of 25%-35% of stocks as Buy and 10%-15% of stocks as Sell; however, the distribution of Buys and Sells in any particular analyst's coverage group may vary as determined by the regional Investment Review Committee. Additionally, each Investment Review Committee manages Regional Conviction lists, which represent investment recommendations focused on the size of the total return potential and/or the likelihood of the realization of the return across their respective areas of coverage. The addition or removal of stocks from such Conviction lists do not represent a change in the analysts' investment rating for such stocks.

Total return potential represents the upside or downside differential between the current share price and the price target, including all paid or anticipated dividends, expected during the time horizon associated with the price target. Price targets are required for all covered stocks. The total return potential, price target and associated time horizon are stated in each report adding or reiterating an Investment List membership.

Coverage groups and views: A list of all stocks in each coverage group is available by primary analyst, stock and coverage group at http://www.gs.com/research/hedge.html. The analyst assigns one of the following coverage views which represents the analyst's investment outlook on the coverage group relative to the group's historical fundamentals and/or valuation. **Attractive (A).** The investment outlook over the following 12 months is favorable relative to the coverage group's historical fundamentals and/or valuation. **Neutral (N).** The investment outlook over the following 12 months is neutral relative to the coverage group's historical fundamentals and/or valuation. **Cautious (C).** The investment outlook over the following 12 months is unfavorable relative to the coverage group's historical fundamentals and/or valuation.

Not Rated (NR). The investment rating and target price have been removed pursuant to Goldman Sachs policy when Goldman Sachs is acting in an advisory capacity in a merger or strategic transaction involving this company and in certain other circumstances. **Rating Suspended (RS).** Goldman Sachs Research has suspended the investment rating and price target for this stock, because there is not a sufficient fundamental basis for determining, or there are legal, regulatory or policy constraints around publishing, an investment rating or target. The previous investment rating and price target, if any, are no longer in effect for this stock and should not be relied upon. **Coverage Suspended (CS).** Goldman Sachs has suspended coverage of this company. **Not Covered (NC).** Goldman Sachs does not cover this company. **Not Available or Not Applicable (NA).** The information is not available for display or is not applicable. **Not Meaningful (NM).** The information is not meaningful and is therefore excluded.

Global product; distributing entities

The Global Investment Research Division of Goldman Sachs produces and distributes research products for clients of Goldman Sachs on a global basis. Analysts based in Goldman Sachs offices around the world produce equity research on industries and companies, and research on macroeconomics, currencies, commodities and portfolio strategy. This research is disseminated in Australia by Goldman Sachs Australia Pty Ltd (ABN 21 006 797 897); in Brazil by Goldman Sachs do Brasil Corretora de Títulos e Valores Mobiliários S.A.; in Canada by either Goldman Sachs Canada Inc. or Goldman, Sachs & Co.; in Hong Kong by Goldman Sachs (Asia) L.L.C.; in India by Goldman Sachs (India) Securities Private Ltd.; in Japan by Goldman Sachs Japan Co., Ltd.; in the Republic of Korea by Goldman Sachs (Asia) L.L.C., Seoul Branch; in New Zealand by Goldman Sachs New Zealand Limited; in Russia by OOO Goldman Sachs; in Singapore by Goldman Sachs (Singapore) Pte. (Company Number: 198602165W); and in the United States of America by Goldman, Sachs & Co. Goldman Sachs International has approved this research in connection with its distribution in the United Kingdom and European Union.

European Union: Goldman Sachs International authorised by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority, has approved this research in connection with its distribution in the European Union and United Kingdom; Goldman Sachs AG and Goldman Sachs International Zweigniederlassung Frankfurt, regulated by the Bundesanstalt für Finanzdienstleistungsaufsicht, may also distribute research in Germany.

General disclosures

This research is for our clients only. Other than disclosures relating to Goldman Sachs, this research is based on current public information that we consider reliable, but we do not represent it is accurate or complete, and it should not be relied on as such. The information, opinions, estimates and forecasts contained herein are as of the date hereof and are subject to change without prior notification. We seek to update our research as appropriate, but various regulations may prevent us from doing so. Other than certain industry reports published on a periodic basis, the large majority of reports are published at irregular intervals as appropriate in the analyst's judgment.

Goldman Sachs conducts a global full-service, integrated investment banking, investment management, and brokerage business. We have investment banking and other business relationships with a substantial percentage of the companies covered by our Global Investment Research Division. Goldman, Sachs & Co., the United States broker dealer, is a member of SIPC (http://www.sipc.org).

Our salespeople, traders, and other professionals may provide oral or written market commentary or trading strategies to our clients and principal trading desks that reflect opinions that are contrary to the opinions expressed in this research. Our asset management area, principal trading desks and investing businesses may make investment decisions that are inconsistent with the recommendations or views expressed in this research.

The analysts named in this report may have from time to time discussed with our clients, including Goldman Sachs salespersons and traders, or may discuss in this report, trading strategies that reference catalysts or events that may have a near-term impact on the market price of the equity securities discussed in this report, which impact may be directionally counter to the analyst's published price target expectations for such stocks. Any such trading strategies are distinct from and do not affect the analyst's fundamental equity rating for such stocks, which rating reflects a stock's return potential relative to its coverage group as described herein.

We and our affiliates, officers, directors, and employees, excluding equity and credit analysts, will from time to time have long or short positions in, act as principal in, and buy or sell, the securities or derivatives, if any, referred to in this research.

The views attributed to third party presenters at Goldman Sachs arranged conferences, including individuals from other parts of Goldman Sachs, do not necessarily reflect those of Global Investment Research and are not an official view of Goldman Sachs.

Any third party referenced herein, including any salespeople, traders and other professionals or members of their household, may have positions in the products mentioned that are inconsistent with the views expressed by analysts named in this report.

This research is not an offer to sell or the solicitation of an offer to buy any security in any jurisdiction where such an offer or solicitation would be illegal. It does not constitute a personal recommendation or take into account the particular investment objectives, financial situations, or needs of individual clients. Clients should consider whether any advice or recommendation in this research is suitable for their particular circumstances and, if appropriate, seek professional advice, including tax advice. The price and value of investments referred to in this research and the income from them may fluctuate. Past performance is not a guide to future performance, future returns are not guaranteed, and a loss of original capital may occur. Fluctuations in exchange rates could have adverse effects on the value or price of, or income derived from, certain investments.