



HOW VALUE CAN TAKE OFF WITH PREDICTIVE AIRCRAFT MAINTENANCE

By Brian Hirshman, Tom Milon, Amanda Brimmer, Ben Brinkopf, Matthew Rabson, and Katherine Smith

DESPITE INTENSE AND LONG-STANDING interest from industry leaders, visions of an AI-enabled future in aviation MRO (maintenance, repair, and overhaul) have been slower to materialize than many expected. And as COVID-19 pushes airlines, OEMs, and MRO organizations toward severe and prolonged cash-conserving measures, it is reasonable to assume that investments in predictive maintenance will be postponed even further.

Our recent experience with players in this space, however, suggests the opposite. While certain “moonshot” (unproven, capital-intensive) predictive maintenance efforts may be less likely, many airlines and MRO organizations, seeking to avoid building back overhead organizations that were sharply reduced as air travel plummeted, are now even more eager to find cost-efficient ways to improve maintenance performance.

Many airlines have already, in recent years, benefited from the use of AI in engine health monitoring (EHM) and have embarked on program optimization cam-

paigns designed to contain costs and increase aircraft reliability. But relatively few have succeeded in early attempts to apply AI to the much more complex component and line maintenance environments. We believe the predictive power of AI could contribute to a step change in airline profitability by removing maintenance-caused disruption from the operation.

The airlines already exploring the potential of predictive maintenance have found their attempts plagued by inadequate sponsorship at the executive level, data fragmentation and usability issues, and myriad downstream implementation challenges.

To overcome these hurdles, we see three critical actions for airlines:

First, our experience with successful implementers of AI suggests that the call for action must come from the executive team, most commonly the COO, and not solely from maintenance leaders. Action needs to be justified by the imperative of keeping planes in the sky, not simply by estimated

cost savings associated with maintenance disruption, which can understate the potential impact. Making predictive maintenance part of an operational excellence push at the airline level gives the effort the strategic consensus and senior influence needed to overcome uncertainty and early frontline resistance. The case for predictive maintenance is the case for operational excellence, fundamental to any airline.

Second, airlines should collaborate more purposefully with one another and with other members of the maintenance ecosystem to move beyond the data-sharing impasse undercutting efforts to build better in-house predictions. By identifying a neutral, incentive-aligned custodian to consolidate and standardize the vast maintenance data set needed to fuel effective prediction algorithms, airlines could get past this debate and accelerate their time to value.

Finally, at the airline level, maintenance needs to interface more effectively with supply chain, operations, and network peers to modify standard work on the frontline, where resistance to change most

often appears. This can introduce timeline and budget risk to the significant investments in process redesign, system changes, and training needed to build and roll out modern predictive tools and programs.

In airlines, as in other industries, the gap between aspirations for AI and the actual impact from the technology can be wide. (See the sidebar.) The challenges are considerable. But we believe the returns from moving to predictive maintenance will justify the investment.

From Engines to Components

Eliminating maintenance-caused disruption requires transforming the maintenance of flight-critical (or “no go”) components from unscheduled to scheduled, no different from the path to value from EHM—prime among AI’s MRO successes to date.

EHM, through prediction and diagnostic tools, enables carriers to avoid unscheduled engine removals by pre-empting in-flight failures and to minimize wasted “green time” of life-limited parts (LLPs) by

BETTING BIG ON AI

Executives in many industries are looking for their company’s first big AI breakthrough. But even in this era of heavy investment, rapid technological change, and rising expectations for AI, many industries remain long on heady visions and short on tangible victories. Some are wondering if the hype around AI has surpassed the technology’s capacity to deliver.

A 2019 BCG GAMMA survey of 2,500 executives from multiple industries shines a light on the gap between organizations’ expectations and results, with seven of ten respondents reporting minimal or no impact from AI to date. And the survey reveals only a tenuous link between investment and impact: whereas 40% of companies making “some” investment in AI have reported

business gains, 60% of companies making “significant” investment report the same. In other words, even full—and fully funded—intentions do not guarantee results.

In any organization considering how to deploy AI, there will be a broad range of proposed use cases, each with distinct cost-benefit profiles and advocates and detractors, all competing for scarce resources. Many organizations are tempted to start with quick wins, which might, as the logic goes, be used to fund future initiatives.

But the survey showed that companies reaping the most benefit from their AI investments tend instead to take on large projects with big potential impact—even if they’re risky.

optimizing shop visit scope. The reliability of these assets has risen in step with the diagnostic capability of condition monitoring algorithms, creating at least the possibility of some types being maintained exclusively on the line and leaving service only when LLP limits require it. With a spare engine costing \$5 million to \$15 million on new platforms, the case for investing in the tools and services to enable this AI breakthrough is straightforward. Consequently, commercializing the technologies to make it possible has become table stakes for engine OEMs.

However, uncovering value in the component environment requires overcoming multiple challenges. To understand these, consider the typical experience in engines: Because an aircraft engine is more or less self-contained, a single party (the OEM) supplies virtually all condition and trend data relevant to a removal decision. The OEM also performs the majority of troubleshooting and repair work. Combining condition and repair data enables an operator to optimize engine maintenance decisions—creating opportunities to remove engines before they fail, extend build standards, and potentially intervene on the line to avoid a removal.

By contrast, decision-relevant data in the component environment of the entire airframe comes from interconnected systems of hundreds of parts, each with varying condition-sensing capabilities. Those systems are manufactured by multiple OEMs, each with its own position on sharing critical data and analysis. Failure in one part may be related to condition deterioration in an associated part, or vice versa. This environment presents airlines with a vastly more convoluted analytical task—more condition readings to parse, a more tenuous correlation between those readings and future failures, a greater need for supplemental data to feed effective predictions, and much more complication in identifying root causes among observable knock-on deterioration in up-line or down-line components. (See Exhibit 1.)

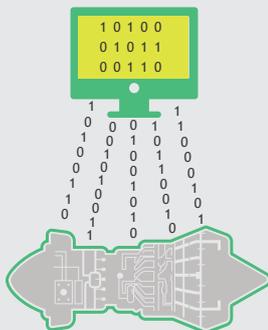
Reducing Disruption

Reducing disruption is perhaps the most-anticipated benefit of predictive maintenance and is almost certain to be its most significant contribution to AI's overall impact on airline operations. Maintenance-caused disruptions cost the industry billions of dollars every year in passenger accommodation costs (room, board, and re-

EXHIBIT 1 | Predictive Maintenance Is Much More Challenging for Airframes Than for Engines

AN ENGINE IS A SELF-CONTAINED UNIT

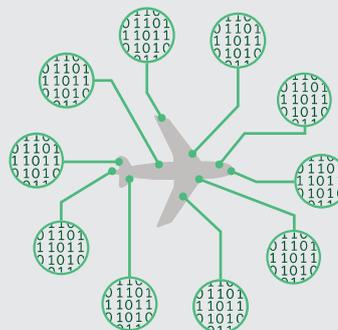
- A single provider may own or manage the entire engine environment of components and can provide a holistic set of condition data.
- The provider can easily track engine health given the single data platform and the interconnectedness of parts.



Source: BCG analysis.
Note: Figure has been simplified.

AN AIRFRAME HAS MANY DIFFERENT INPUTS

- Multiple providers, that may not communicate or share data with one another, supply various airframe components.
- Operators must combine varying condition readings, gather supplemental data, and attempt to analyze this fragmented data to develop accurate predictions.



System examples: hydraulic, pneumatic, APUs, etc.

booking), additional gates, crew compensation, wasted fuel, and aircraft recovery. And beyond these costs lie the less quantifiable but more detrimental impacts of lost lift, lower productivity, and erosion of passenger loyalty. Even marginal improvement in this area is expected to unlock material value for operators. The importance of on-time performance is dogma in any airline, and these metrics are obsessively tracked for good reason.

Much of what's been written about predictive maintenance in aviation describes a glossy vision of the future, perhaps detached from the complicated realities of daily operations. In this vision, a component failure warning from a sensor-enabled aircraft reaches the operator's maintenance operations center (MOC) midflight, setting off a symphony of tightly coordinated activity and resulting in a serviceable replacement in the hands of a waiting mechanic as the aircraft touches down. Tada!

This oversimplified view ignores the complexity involved both in making reliable predictions and in integrating those predictions into systems able to define and enact disruption-preventing maintenance activities. While a continual feed of condition data is the key originating feature of predictive maintenance, and its production can be taken as a given for major systems on modern aircraft, it takes more—much more—than smart sensors and a live data feed to make predictive maintenance real.

Generating value from predictive maintenance depends on overcoming three challenges:

DATA VOLUME AND ACCESS

Far from the “magic wand” vision of the future often described, carriers exploring their opportunity in predictive maintenance often discover condition data from their own fleet to be insufficient to inform credible predictions. And for a simple reason: for a majority of parts, particularly high-value parts, there are not enough failures within a single fleet to validate those predictions. One airline testing a failure algorithm across a set of hundreds of compo-

nents, for example, was successful only in validating predictions for coffee makers.

Some operators choose to subscribe to manufacturer alert services as a basis for their predictive maintenance program. Though these services provide abundant condition alerts that often do prevent a delay-causing failure, some users have discovered that acting on the recommendations spikes maintenance costs. In these cases, false positives drive up removal rates, generating no-fault-found events (in which a part is inspected and found to be serviceable), extending overnight ground time and demanding additional inventory to support higher removal frequencies.

Ultimately, condition data in isolation is rarely enough to generate failure predictions consistently accurate enough to both drive down disruption and contain maintenance costs. Though carriers with large, uniform fleets and better data sets may see more success, the homogeneity of condition data hinders progress for most. To surmount the issue, operators need access to a greater and more diverse stream of condition, trend, and failure data—one with a higher count of failures, a greater diversity of failure modes, and a record of corrective actions taken. This is the fuel needed for effective prediction modeling on low-removal, flight-critical parts.

To build such a data set, airlines have to resolve a years-old stalemate over ownership of operation, condition, maintenance, and reliability data. Though airframe and component OEMs record, store, and analyze part-level condition and removal data across a massive number of aircraft, they are prevented by confidentiality agreements from sharing it across airlines in a granular enough form to bring diagnostic utility to predictive algorithms. Several players have attempted to consolidate and anonymize multicarrier condition, failure, and maintenance data into a single pipeline, but, so far, few airlines have shown a willingness to entrust their data with parties that are perhaps seeking to commercialize it, despite assurances of confidentiality and authorized use.

Ongoing efforts to broaden and bring order to a collective data set are leading to fragmentation, with no clear frontrunner emerging or signs of willingness among large numbers of airlines to participate. To resolve the impasse, airlines should consider pushing harder for a common data pipeline built in their own image.

Such a pipeline would need to achieve three primary aims:

- **Anonymity and Confidentiality.** This is the first source of resistance when airline maintenance and IT leaders consider contributing to a common data pipeline, consistent with a history among airlines of holding tight to maintenance and reliability data despite frequently collaborating in other technical support areas. Potential participants will insist their data be faithfully protected both technically and legally, and properly aggregated and sanitized before being made available for user query.
- **Commensurability.** Operators of large fleets may view participation in a shared data pipeline as asymmetrical—a proposition in which the value of their data contribution outweighs the potential benefit. And the non-participation of large carriers may dissuade next-largest operators from contributing, and so on. Large carriers, though potentially contributing in an outsized way, also stand to benefit from incremental improvement to predictions by virtue of their scale and complexity, which make disruption particularly costly, at least on an absolute scale. And were the pipeline to take shape as a commercial endeavor, net data contributors might eventually find it a source of ancillary revenue.
- **Neutrality.** Airline cooperation also depends on the perceived trustworthiness of the pipeline’s custodian. Airlines have thus far largely resisted third-party efforts to develop and grow a data pipeline, often for this reason. A custodian with a fiduciary duty to,

rather than a commercial interest in, the pipeline’s contributors may be the structure necessary to address this concern.

While IATA or a new joint venture may be viable options, in our view, SITA—member-owned, with a mandate to represent the global aviation community—is perhaps the most credible candidate for this custodial role. SITA already has a fiduciary relationship to members, the ability to store, manage, protect, and transmit massive streams of data, and experience as a paracommercial enterprise. And SITA is already experimenting with adjacent offerings that might find appealing white space in predictive maintenance.

SITA launched the MRO Blockchain Alliance in 2019 and unveiled a plan to bring industry stakeholders (MROs, airlines, lessors, logistics providers) together to set a global standard for the use of blockchain technology to record and track aircraft parts, possibly including repair data, which is notoriously difficult to aggregate and operationalize. SITA, through its SITAONAIR subsidiary, has also rolled out the e-Aircraft data hub to collect unstructured real-time data from aircraft operations and feed it to OEMs so it can be analyzed and built into digital services and applied to equipment modifications and maintenance regimes.

One can envision a future in which the condition data SITA’s e-Aircraft collects and collates for OEMs is made available directly to airlines, supplemented by repair histories digitized by the MRO Blockchain Alliance, transmitted through largely existing infrastructure and encrypted by blockchain technology. At critical mass, this combination would provide the incentives needed to promote contribution to and use of a shared data pipeline:

- The value of combined condition, failure, and diagnostic data at scale would vastly increase prediction quality. With a robust and evergreen data set, airlines could generate predictions for more parts with more precision, avoid false positives, and inform the develop-

ment of preemptive maintenance tasks to increase reliability.

- Centralizing and safeguarding data through an impartial body, with a mandate to serve its contributors and users, would address the primary concern among would-be data contributors—the risk of data misappropriation.
- Transmitting the data through existing infrastructure on a standard protocol would simplify the task of access and collation for an airline, whose IT resources are frequently stretched, particularly in maintenance.

With proofs of concept on current initiatives, SITA—or another party—may provide potential data contributors with the value and comfort necessary to lobby for a global pipeline. And with critical mass, operators hamstrung in their efforts to develop predictive maintenance capabilities may finally have the analytical grist they need.

OPERATIONAL INTEGRATION

Airlines have begun to incorporate on-board component condition readings into AI-enabled prediction tools but, in most cases, have not integrated those tools with operations. Until they do, they remain saddled with a “brute force” approach that deprives them of the operational agility needed to take advantage of improving prediction technology, the value of which can be reaped only if preventive action is taken inside a prediction’s window.

Particularly in the early stages of prediction development, as time to the predicted moment of a part’s failure increases, precision declines (the confidence interval widens) and utility rises (sufficiency of the time available to pre-empt). Shortening post-prediction response time, therefore, addresses both sides of the risk that predictive maintenance attempts to balance:

- **Risk of late removal.** If a carrier is unable to react within the prediction window, the part will fail, potentially causing a disruption event.

- **Risk of early removal.** If a carrier prioritizes preventive response, it must utilize less-precise “far-out” predictions, introducing the risk of premature removal.

As AI-enabled prediction quality improves over time, an airline’s failure to develop the infrastructure needed to turn predictions into action-enabling commands, rather than the accuracy of the prediction itself, will become the primary obstacle to creating value. As is, downstream response has not evolved in step with the shrinking time horizon of accurate predictions, and airlines are often not nimble enough to prevent the predicted disruption.

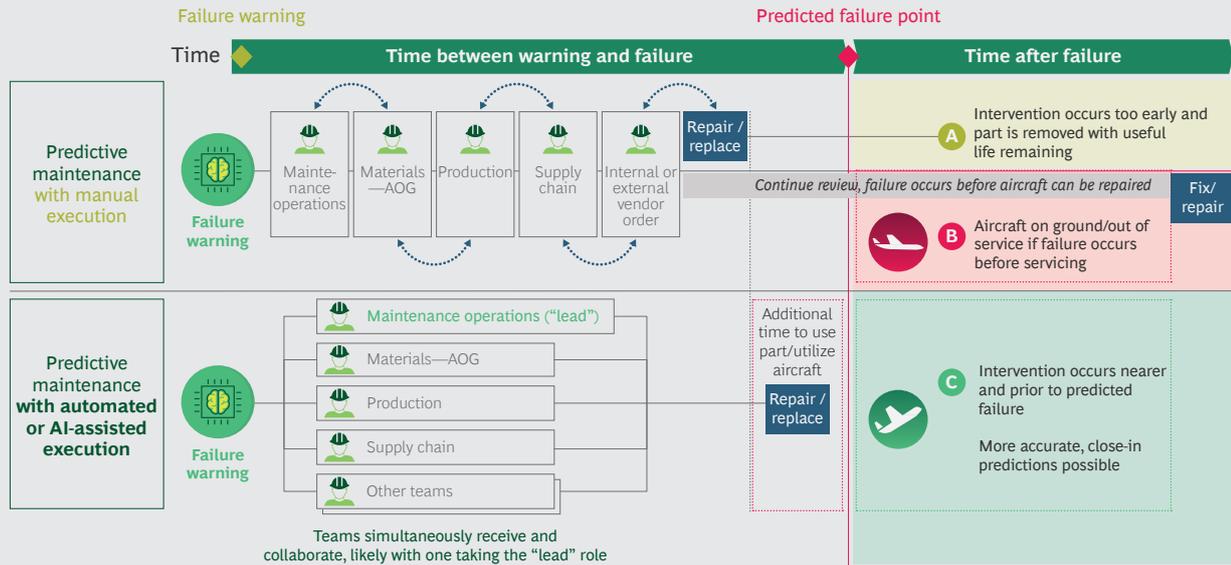
To prevent disruptions, downstream processes need to respond to failure predictions simultaneously rather than as serialized steps that depend on manual work and build in waiting time. (See Exhibit 2.) Tightly integrating these processes with the evolving prediction tool is, unavoidably, a significant undertaking—both in time and investment—involving changes to people’s work routines and legacy IT systems. See Exhibit 3 for an overview of the current and desired state of predictive maintenance.

INTELLIGENCE TIED TO STRATEGY

A basic premise of predictive maintenance is the continuous balancing of late removal risk, and its potential impact on operational disruption, and early removal risk, and its potential cost in no-fault-found events and lost green time. Whereas the hard costs of early removals can be estimated with some confidence, assigning value to avoiding disruption is full of complex dependencies and assumptions.

At many airlines, an investment case premised on disruption avoidance can meet with dubious approvers more apt to prioritize projects justified by hard dollars and known return horizons. This early hurdle can deprive a maintenance-led predictive maintenance initiative of the “venture capital” and senior consensus needed just to get started.

EXHIBIT 2 | Significant Value in Pairing Predictive Maintenance with Automated or AI-Assisted Execution



Source: BCG analysis.

EXHIBIT 3 | Where We Are—and Where We Need to Be—with Predictive MRO in Aviation

	CURRENT STATE	NEEDED STATE
Engineering validation	<ul style="list-style-type: none"> Predictions validated or rejected by an engineer through manual compilation, review, and analysis of corroborating data 	<ul style="list-style-type: none"> Predictions routed to engineering in real time and queued by urgency rather than first-in, first-out Validation largely automated by assembling historical removal and repair data on a serial- and part-number basis
Mid-term planning	<ul style="list-style-type: none"> Maintenance planner queues validated removal prediction into planned maintenance Planner annually generates a task to replace the affected component 	<ul style="list-style-type: none"> Validated removal prediction assigned to an optimized RON date and location based on operating schedule, manpower capacity, material availability, and station capability
Material requisition	<ul style="list-style-type: none"> Supply chain generates a tools and materials forecast based on queue of short-term planned maintenance and queries the inventory system 	<ul style="list-style-type: none"> Material and tooling requirements defined in the authorized task are pushed to the inventory system In-stock materials are flagged for shipment to the recommended maintenance location, with approval of logistics team
Procurement	<ul style="list-style-type: none"> Depending on availability, supply chain issues a PO for the needed tools and materials 	<ul style="list-style-type: none"> Material and tooling requirements defined in the authorized task are pushed to the procurement system Out-of-stock material with required delivery-by date is ordered and shipped from approved vendors
Short-term planning	<ul style="list-style-type: none"> A short-term maintenance planner manually monitors the prediction deadline and, with OCC, assigns task to a maintenance location 	<ul style="list-style-type: none"> Prediction status is visible to a short-term maintenance planner with changes in risk highlighted in an alert system Planner intervenes only to reassign the associated task to a date and maintenance location if needed
Aircraft routing	<ul style="list-style-type: none"> OCC operator routes the aircraft to the chosen maintenance location in advance, manually coordinates 	<ul style="list-style-type: none"> Additional time and automatic routing allow aircraft to fly more of its schedule and have part replaced at scheduled maintenance location
Material handling	<ul style="list-style-type: none"> Supply chain orders tools and materials and sends to the chosen maintenance location, if necessary 	<ul style="list-style-type: none"> Supplies are ordered based on collective team decision Materials are automatically routed to aircraft based on revised schedule
Maintenance provider	<ul style="list-style-type: none"> Receives aircraft with tight turnaround time and limited context on repair needs 	<ul style="list-style-type: none"> Full case file provided in advance with parts and tools prepared

Source: BCG analysis.

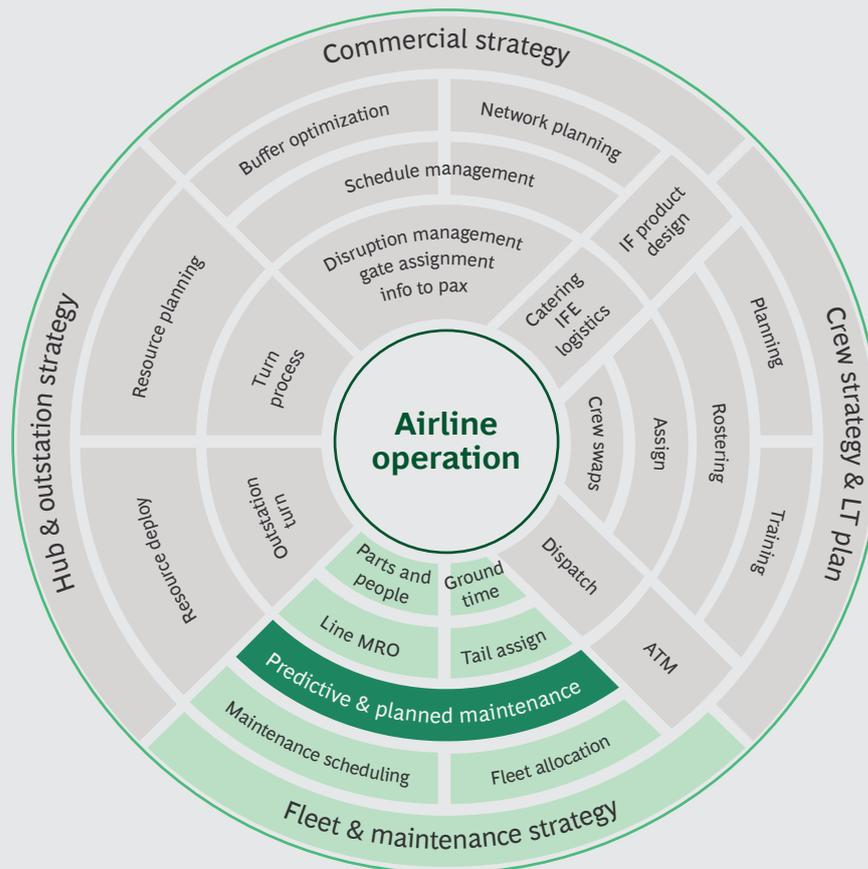
But the repercussions of maintenance disruption stretch beyond the KPIs of the maintenance function to adjacent areas that must collaborate to deliver flawless operations—flight crew scheduling, ground crew scheduling, network planning, ground operations, and more. (See Exhibit 4.)

It follows, then, that getting predictive maintenance off the ground should not fall solely to maintenance leaders. At many carriers, maintenance and engineering is physically isolated in an airport facility instead of embedded at headquarters, or is functionally siloed as a support organization instead of included in senior decision bodies. This distance from the airline’s core can undercut efforts to attract the

funding, interest, and urgency needed to launch and keep momentum behind an ambitious initiative.

Airlines making progress in predictive maintenance, in our experience, differ in their upfront approach, relative to a typical adopter, in two critical areas: First, they encompass their predictive maintenance program in an airline-wide effort to drive operational excellence, through AI or otherwise. Second, they assign one of the company’s most influential leaders, most often the COO, to sponsor all aspects—defining success metrics and milestones, assigning internal and external resources, and ensuring regular visibility at the executive level.

EXHIBIT 4 | Predictive Maintenance Can Improve the Airline’s Overall Operation



Predictive maintenance is an integral part of an airline’s operation and can unlock value

Source: BCG analysis.

One airline that has achieved step-change success in reducing maintenance disruption took this approach, agreeing to relieve the maintenance organization of the onerous task of proving financial payback, instead asking only for a tally of avoided disruptions. As long as credible progress could be demonstrated, financial and managerial support flowed to the project.

With this foundation, maintenance decision intelligence (which follows a prediction) can be suffused with an airline's full range of inputs and used to optimize the operational response around the airline's strategy, not a maintenance KPI. Whereas maintenance-controlled decision intelligence may solve for, say, technical dispatch reliability or AOG duration, a strategically aligned decision logic may advise differently in many, or even most, cases.

For example, consider some scenarios leading up to the pending failure of an expensive flight-critical component, considered likely to cause the cancellation of a full flight on a leg from the airline's busiest hub:

- **Airline A, prioritizing customer loyalty**, may overweight the satisfaction of high-status passengers and choose to purchase a serviceable component, at spot price, for installation at the first line stop with sufficient ground time to prevent any delays and the potential cancellation.
- **Airline B, prioritizing employee satisfaction**, may stretch to preempt crew disruption and choose to source an exchange, rather than replacement, component, at spot price, for installation at the line stop, relieving crew of an unplanned overnight but allowing for potential flight delays.
- **Airline C, minimizing direct cost**, may choose to wait to replace the component until the aircraft rotates to the hub, converting the potential cancellation to a long maintenance delay and avoiding any investment in an additional component.

By contrast, leaving the post-prediction decision to the siloed logic of the maintenance group may, for example, allow technical dispatch reliability or AOG avoidance to be maximized at great cost—consistent with the objectives of Airline A but not with those of Airline C in many instances. Strategically aligning that decision logic demands close involvement from cross-functional leadership before a single line of code is written. Developing the universe of adverse scenarios that AI ultimately attempts to solve for is a human-led, data-aided endeavor. And assigning values to those scenarios, similarly, can be quantified only to a degree before the judgment of the leadership team must generate prioritizing (and subjective) coefficients to weight decision rules reflecting the airline's strategic intentions.

The Path to Flawless Operations

The challenges are considerable, but there is a prize worth pursuing in predictive maintenance. The classic sources of value surmised by any industrial player investing in predictive maintenance—increasing material and labor efficiency, reducing inventory investment and, most critically, minimizing asset downtime—can all be tapped using the tools, technology, data, and expertise becoming more broadly available.

To get to value, airlines first need to treat maintenance as one critical contributor to operational excellence among many interconnected contributors, view an effective predictive maintenance capability as the path to eliminating associated disruption, and prioritize resources and investments accordingly. With strategic buy-in, carriers may choose to push harder to improve the fragmented data environment that has hamstrung efforts to ply algorithms with the robust repair and removal history needed to develop consistently effective predictions. Once this availability challenge can be met, airlines will have to invest aggressively to protect, standardize, and embed that data in tools that interface with the airline operation—where the step-change benefits of predictive maintenance really lie.

Our research shows that companies getting the most from AI are the ones that place the biggest bets. At airlines, betting big on AI means betting on the value of aircraft up-time and a punctual, predictable opera-

tion. Predictive maintenance is one of many AI-enabled levers that must be pulled, at the right time and in the right way, to deliver flawless operations.

About the Authors

Brian Hirshman is a managing director and partner in the Dallas office of Boston Consulting Group and an expert in aviation and aerospace. You may contact him by email at hirshman.brian@bcg.com.

Tom Milon is a partner in BCG's Philadelphia office. You may contact him by email at milon.tom@bcg.com.

Amanda Brimmer is a managing director and partner in the firm's Chicago office and a core member of the Industrial Goods practice. You may contact her by email at brimmer.amanda@bcg.com.

Ben Brinkopf is a principal in BCG's Dallas office. You may contact him by email at brinkopf.ben@bcg.com.

Matthew Rabson is a principal in the firm's London office. You may contact him by email at rabson.matthew@bcg.com.

Katherine Smith is a principal in BCG's Atlanta office. You may contact her by email at smith.katherine@bcg.com.

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we help clients with total transformation—inspiring complex change, enabling organizations to grow, building competitive advantage, and driving bottom-line impact.

To succeed, organizations must blend digital and human capabilities. Our diverse, global teams bring deep industry and functional expertise and a range of perspectives to spark change. BCG delivers solutions through leading-edge management consulting along with technology and design, corporate and digital ventures—and business purpose. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, generating results that allow our clients to thrive.

© Boston Consulting Group 2020. All rights reserved. 8/20

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on Facebook and Twitter.