WHITE PAPER

# The Existence of a Positive-Sum Game Among Airlines

COVID Task Force

# Preface

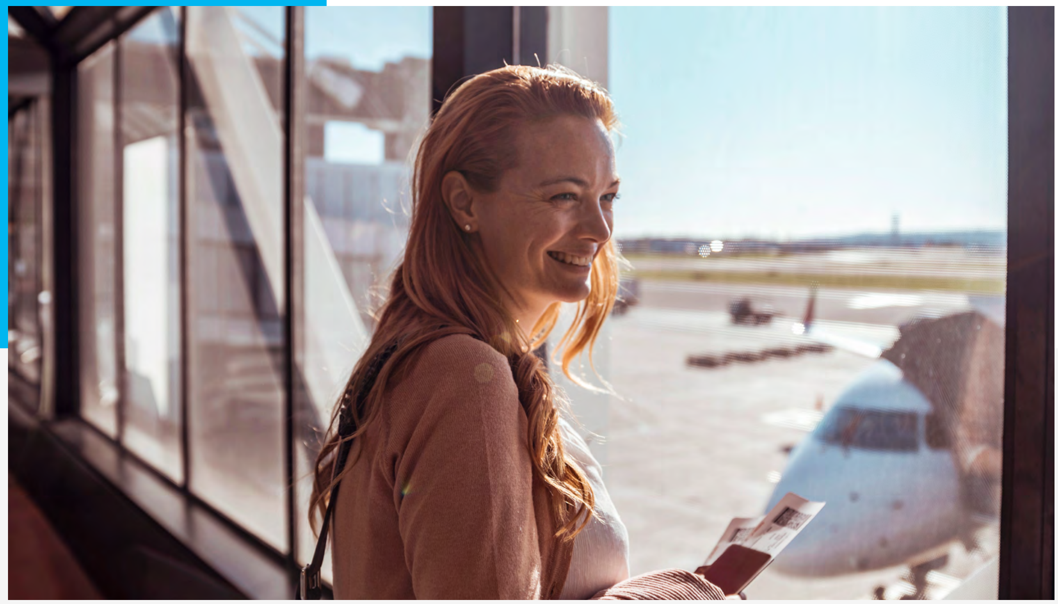

This whitepaper summarizes the work of the COVID Task Force (CTF) that started in March 2020. Since many webinars, conference submissions, and a series of eight blog posts have already been created to document our progress and findings along the way, we will focus on topics that have not been discussed elsewhere and summarize materials that have already been published. Moreover, due to the overwhelming number of airline participants in the CTF compared to our non-airline participants, this whitepaper will place more emphasis on our work with the airlines.

# Introduction

Contrary to common beliefs, competition and cooperation are not opposites of each other. They co-exist and are inherent to every social structure across all cultures around the world.

Although the global economy is largely driven by competition for scarce resources among agrarian societies in the pre-industrial age, the innovations of the post-industrial era enable us to use our limited resources so efficiently that we create a huge surplus. Consequently, we specialized, became more co-dependent, and learned to collaborate to innovate further, which creates a positive feedback loop.

Historically, we've seen that it's possible to turn a zero-sum game into a positive-sum game through collaboration and innovation. This made me wonder whether there is any room for cooperation in the airline industry where competitive pressure is extreme and profit margin is razor-thin. Can we innovate together to grow the pie, so everyone gets more pie?

Under normal business conditions, this would be an interesting thought experiment, a kind of intellectual exercise limited to the domain of academics. It would be extremely difficult to test and quantify this hypothesis in reality. However, we've all experienced the crisis of the century in 2020. Despite the detrimental impact of COVID-19, it's also an opportunity of a lifetime for the academician in me. Not only can we ask the "what ifs," but if we can design a sort of "natural experiment" just right, we could potentially measure the benefit and value of a large-scale collaboration among the airlines.

# The Perfect Storm and Our Call to Action

How can we drive a large-scale collaboration among the airlines? Under normal business operations, this would never happen. But the onset of the COVID pandemic is a perfect storm in the brewing.

First, the airlines need help. Because consumer travel behavior is severely disrupted due to the various containment measures across the globe, future travel patterns are much harder to forecast. Without an accurate estimate of future demand, airlines can't use revenue management (RM) to effectively optimize the price for their seats, which puts the airlines' future profitability at risk.

Second, we, PROS, can help. Our core competency in RM and demand forecasting had been serving the airline industry for many decades. In times of uncertainty, our deep data science expertise is the greatest value we can contribute to the industry. With the right data, we can help businesses get a clearer glimpse of the future through our machine learning (ML) models. Moreover, we are in a unique position to help because we serve a majority of the global airlines.

Finally, the most challenging part of the abovementioned investigation is to actually get a substantial number of airlines to participate in such a collaboration. Obviously, we cannot coerce any business to cooperate as we are all operating under a free economy. Moreover, since we value transparency in the way we conduct our business, including research, we cannot conduct this investigation using any blind or double-blind methodologies. The only remaining option is to have the airlines opt into this collaborative effort voluntarily. This is next to impossible because so many aspects of the airlines' business have always been treated as a zero-sum game.

Clearly, this will require some investments. We must invest to create substantial value to the airlines, that is above and beyond their immediate needs to respond to the pandemic. After several virtual office hours with the airlines to understand their needs, we formulated our call to action and assembled the COVID-19 Task Force (CTF) with members from our elite data scientists team. Born out of the strong belief that we will all benefit from an industry-wide collaboration, the CTF is committed to helping our customers navigate through this unprecedented crisis.

# The Purpose of the COVID Task Force

**The CTF was set out to address the most frequently asked inquiries from our customers, which were collected from the initial office hours. They can be grouped into the following three categories:**

1 Managing the Immediate Crisis

2 Forecasting Business Recovery

3 Resetting to The New Normal

| Blog Posts |
|---|
| PROS Assembles a COVID-19 Task Force to Help Customers |
| COVID Task Force 1: Managing Your Business Under Crisis |
| COVID Task Force 2: Revenue Management Under Lockdown |
| COVID Task Force 3: What Are We Doing With Your Data |
| COVID Task Force 4: Building A Crystal Ball for the Airlines |
| COVID Task Force 5: How Airlines Will Return to Cruising Altitude |
| COVID Task Force 6: From Prediction to Prescriptive Actions |
| COVID Task Force 7: The Future of Revenue Management |
| **Conference Submissions** |
| AGIFORS RM Study Group 2020 |
| Rice Data Science Conference 2020 |
| Outperform 2020 |

**TABLE 1**

The category 1 inquiries are urgent. At this point, we have been actively working with our customers and created a multitude of resources to help them deal with the crisis. A growing list of publicly accessible resources can be found in Table 1.

Addressing the remaining inquiries in categories 2 and 3 will be the bulk of the CTF's work. This essentially involves building a large-scale machine learning (ML) model that forecasts the future booking trend. Using this forecasted global demand trend, the CTF will conduct a customized "science study" using customer-specific data to re-optimize our RM system for each of the participating airlines. This will ensure our RM system works properly under the new normal where the demand pattern is significantly changed.

Although we will deliver all these at no cost to the participants in our CTF study, we are going above and beyond to provide extra perks for those who chose to participate.

1. First, participants can provide inputs to augment our standard forecast, making this a customer-specific forecast.
2. Participants also receive more granular and customer-specific predictions based on customer augmented input.
3. Finally, we will also offer time to help participants make sense of and strategize the use of our deliverables, including webinars and individual consultation sessions.

Although our optimized carrier-specific recommendations can be implemented as changes to the drift parameters and censoring patterns, it's up to the participants whether they want to implement these changes in their respective RM systems.

# Our Partners: The Airline Participants and Beyond

To develop such a planetary-scale booking forecast model, we need a huge amount of data. Since no single carrier covers the entire globe, this is where an industry-wide collaboration would be extremely helpful. If every airline contributes their booking data to train our booking forecast model, our model will be more predictive and more accurate. And if our forecast can more accurately reflect real demand, everyone gets the benefit from tuning their RM to operate more optimally under the new normal.

Although this line of argument is logically intuitive and intellectually pleasing. It does, however, require a leap of faith from the airlines. They would have to contribute their booking data to be combined with others in an aggregated and anonymous fashion to train our global booking forecast model. And this, to my knowledge, is something that airlines have never done before. Since an airline's market share is secret competitive intelligence, it's understandable that most airlines would be reluctant to share this data. However, in dire times, we humans can often look beyond the competition in front of us and see the bigger picture.

When the CTF was assembled, the call to action was an open invitation for all PROS customers. Due to airlines' extreme sensitivity around sharing their booking data, we initially expected more B2B customers to participate than airlines. However, the response from the airline industry was overwhelming compared to our B2B customers. That is why this whitepaper is airline-focused. Nevertheless, we will cover some of the relevant works with B2B participants at the end.

Although some airlines required a little bit of convincing, we did end up with roughly two dozen global airlines contributing their booking data to our CTF effort. By doing so, they are contributing to the entire industry by making the booking recovery forecast more accurate and more generalizable for all participants. We are honored to facilitate this historic collaboration with our airline partners, as this is truly unprecedented.

# Data and Features in our Global Demand Forecasting Model

With all the booking data from our CTF participants, we are now ready to develop the global demand predictor. However, COVID is a black swan that has never happened before, so its progression and impact cannot be accurately predicted using history alone. Therefore, we must use third-party data sources to augment our model. To improve our model, we've been gradually incorporating more and more third-party data into our model over the past 20+ months.

We started with just epidemic data and government lockdown data as third-party data predictors. And the latest version of our model is ingesting 5 categories of data from the different providers (Table 2).

| Data Categories | Data Providers |
|---|---|
| Epidemic data | Center for System Science and Engineering (CSSE) at Johns Hopkins Univ |
| Government response data | Blavatnik School of Government at Univ. of Oxford |
| COVID-19 epidemic projections<br>Cell phone mobility projections | Institute for Health Metrics and Evaluation (IHME) at UW Medicine |
| Vaccination data | Our World in Data by Global Change Data Lab (GCDL)<br>in partner with Oxford Martin School at Univ. of Oxford |

**TABLE 2**

In addition to leveraging more data sources, we also engineered more informative features from these raw data. One of the most predictive features we engineered is the "flight route openness" feature (Figure 1). It is derived from the government response data's stringency index, $r(t)$, which is a score between 0 and 100% that varies over time, $t$. Since greater stringency implies the country is locked down, the openness of the country, $c$, for air travel can be represented as $[1-r_c(t)]$. Moreover, since both the origin country, $o$, and destination country, $d$, must be open for a flight to operate on the route between o and d, this route's openness is defined as $R_{od}(t) \triangleq [1-r_o(t)]\,[1-r_d(t)]$.

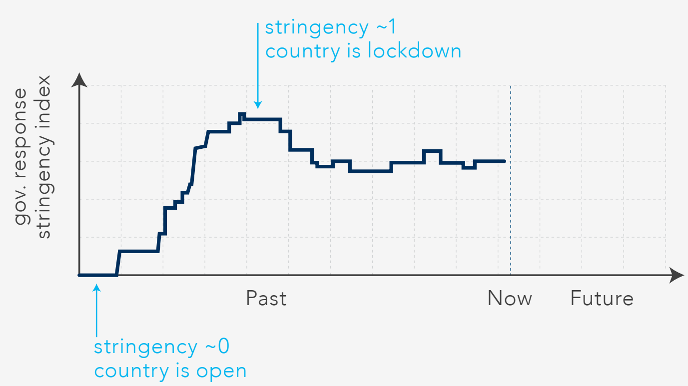


$$R_{od}(t) = [1-r_o(t)] \times [1-r_d(t)]$$

Flight route (o&d)'s openness to travel | Origin's openness to travel | Destination's openness to travel

**FIGURE 1**

| Epidemic Features |
|---|
| New Case/100k at Origin & Destination |
| New Death/100k at Origin & Destination |
| Cumulative Case/100k at Origin & Destination |
| Cumulative Death/100k at Origin & Destination |
| Mortality Rate at Origin & Destination |
| **Government Response Features** |
| Flight Route Openness |
| **Vaccination Features** |
| Vaccination Rate at Origin & Destination |
| **Booking Related Features** |
| Carrier |
| Fare Compartment |
| Origin & Destination Region |
| Days Prior |
| Flight Type (domestic, intra- or inter-region) |
| Bokings (2019, 1 week ago & 2 weeks ago) |
| Cancels (2019, 1 week ago & 2 weeks ago) |
| Net Bookings (2019, 1 week ago & 2 weeks ago) |

**TABLE 3**

There are also many simple features we derived from the ingested raw data throughout the course of our investigation. The feature set used in the latest iteration of our model is summarized in Table 3.

# Training and Optimizing Our Model

As with many machine learning exercises, we begin our modeling journey using highly interpretable linear models as a baseline. Then move on to more complex, but less interpretable models, such as elastic net and random forest. Due to the significant difference in predictive performance, we have settled on using random forest as our model of choice.

The predictive power of the random forest model is not without cost. Because random forest is a regression tree-based ensembled model, it has many hyperparameters that need to be tuned and optimized. These model hyperparameters are set before the training process, so they are not determined by the training data. For example, the number of trees, maximum tree depth, subsampling rate, number of allowed features in each tree, etc.

Fortunately, the model's performance usually does not vary significantly with these hyperparameters, so they only need to be optimized once for each model. However, when the structure or input features of the model change, the hyperparameters need to be reoptimized. This is a computationally-intensive process that takes 126 hours on a cluster of 8 nodes, each with 8 cores and 28 GB of memories.

# From Prediction to Forecasting: Our Unique Approach

If you are a cautious reader, you might have noticed that we have not yet discussed how we are using the epidemic projection and mobility projection data from IHME. To understand how we leverage these data from IHME, we must understand the difference between "prediction" and "forecast." Despite the similar connotation of the English words, they are two different math problems.

We have trained and optimized our random forest model so it's able to accurately predict bookings from epidemic data, government response, and vaccination rate. That means we have a good understanding of how bookings will change as the epidemic, government response, or vaccination rate changes. For example, if a country lifted a lockdown today, our random forest model could tell us how bookings will increase today.

Although this prediction is informative, it's not very actionable, because there is no lead time for the carrier to prepare for this increase in bookings. To make this booking prediction more actionable, we must know the increased booking far enough in advance so the airline can act and prepare for it. This is where we need forecasting. Forecasting is generally a harder problem, because we must push our predictions forward in time.

In other words, we must predict bookings in the future, say tomorrow, or a month from today. This means we will need future data for the epidemic, government response, and vaccination rate (tomorrow or a month from today), which we will not have today, yet.

Since no one has future data, the future values for the epidemic, government response, and vaccination rate must be forecasted as they are the required input to our random forest model. These required inputs will appear to limit the usability of our model, since our model can only be used when you have all the inputs. However, the required inputs allow our model to more accurately reflect reality and make our prediction more contextual. Moreover, it allows our model to be augmented by our CTF participants. Local knowledge about the epidemic, government response, and vaccination can be incorporated from each participating airline via the model's input.

However, it seems that we are simply turning one hard problem (i.e. forecasting bookings) into another equally hard problem (i.e. forecasting the epidemic, government response, and vaccination rate). This is where IHME's data comes in handy. IHME provides a rough estimate of the future values for the epidemic and vaccination. We preprocess these projections by bucketing and normalizing them before using them as inputs to our global booking predictor.

Moreover, IHME provides a future projection of mobility at the country level using cell phone movement data. Since many government-implemented containment measures restrict people's movement, a government's response stringency is strongly anti-correlated to the mobility of its constituents. Hence, IHME's forecasted future mobility data can serve as a good proxy for future government responses. Now, we have all the ingredients to turn our global booking predictor into a global booking forecaster.

Although we are indeed turning one hard problem into another hard problem, we are allocating the best experts to deal with the problem they are best at solving. We are letting the medical experts and the epidemiologists solve the problem that they are most equipped to solve (i.e. forecasting the epidemic and government response). And we are letting PROS scientists solve the problem that we are best at solving (i.e. forecasting bookings). This convergence of diverse expertise is what makes our model unique.

# Our Last Booking Forecast in Nov. 2020

Now that we have a good model, we can use it to make a forecast about how booking will recover in the future. We have already published an initial regional level forecast in [Business Travel News] in Dec. 2020. A more detailed discussion of this recovery forecast can be found in our [CTF blog update #5].

Keep in mind that this booking recovery projection was made using data up to Nov. 2020 for a future forecast window 3 to 6 months out (from Nov. 2020). That means the forecasted booking level was for the window from Feb. to May of 2021. Since this forecast window has already passed, we can compare our forecast to the observed booking level to assess how we did (see Figures 2, 3, and 4).

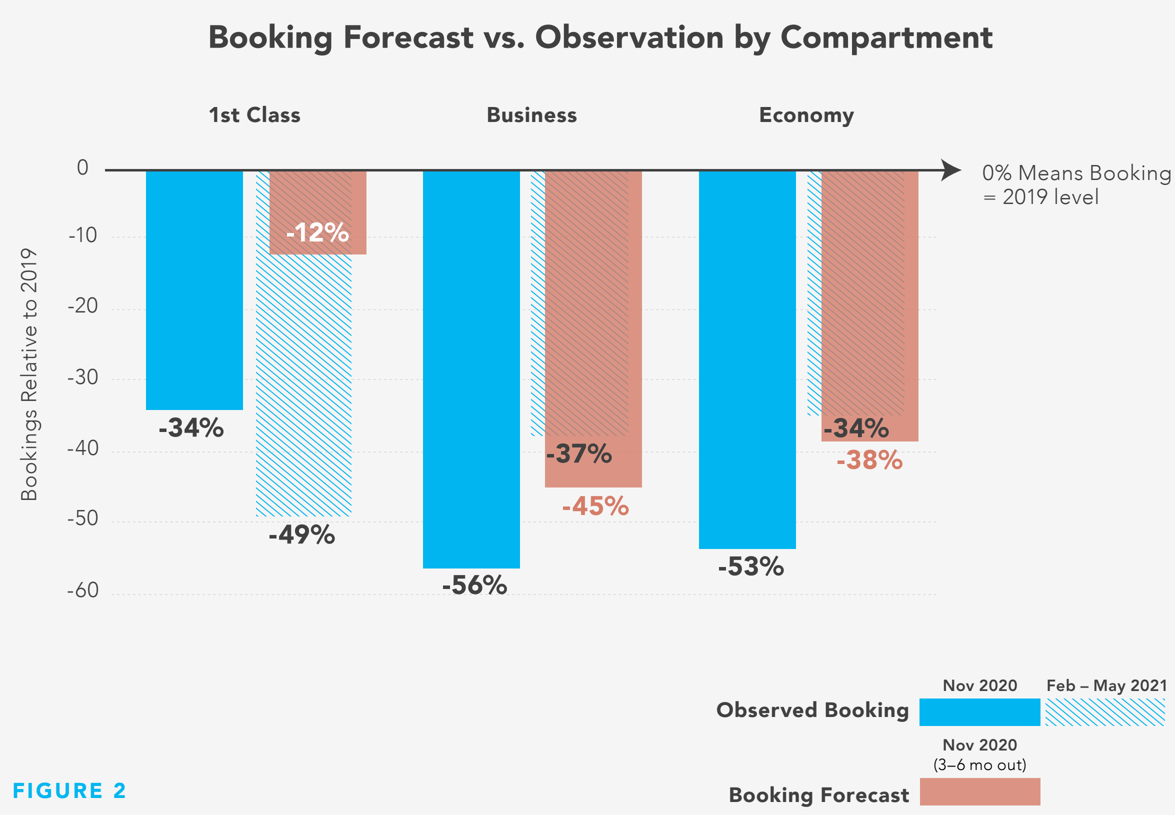


FIGURE 2

## Booking Forecast vs. Observation by Region and Flight Type

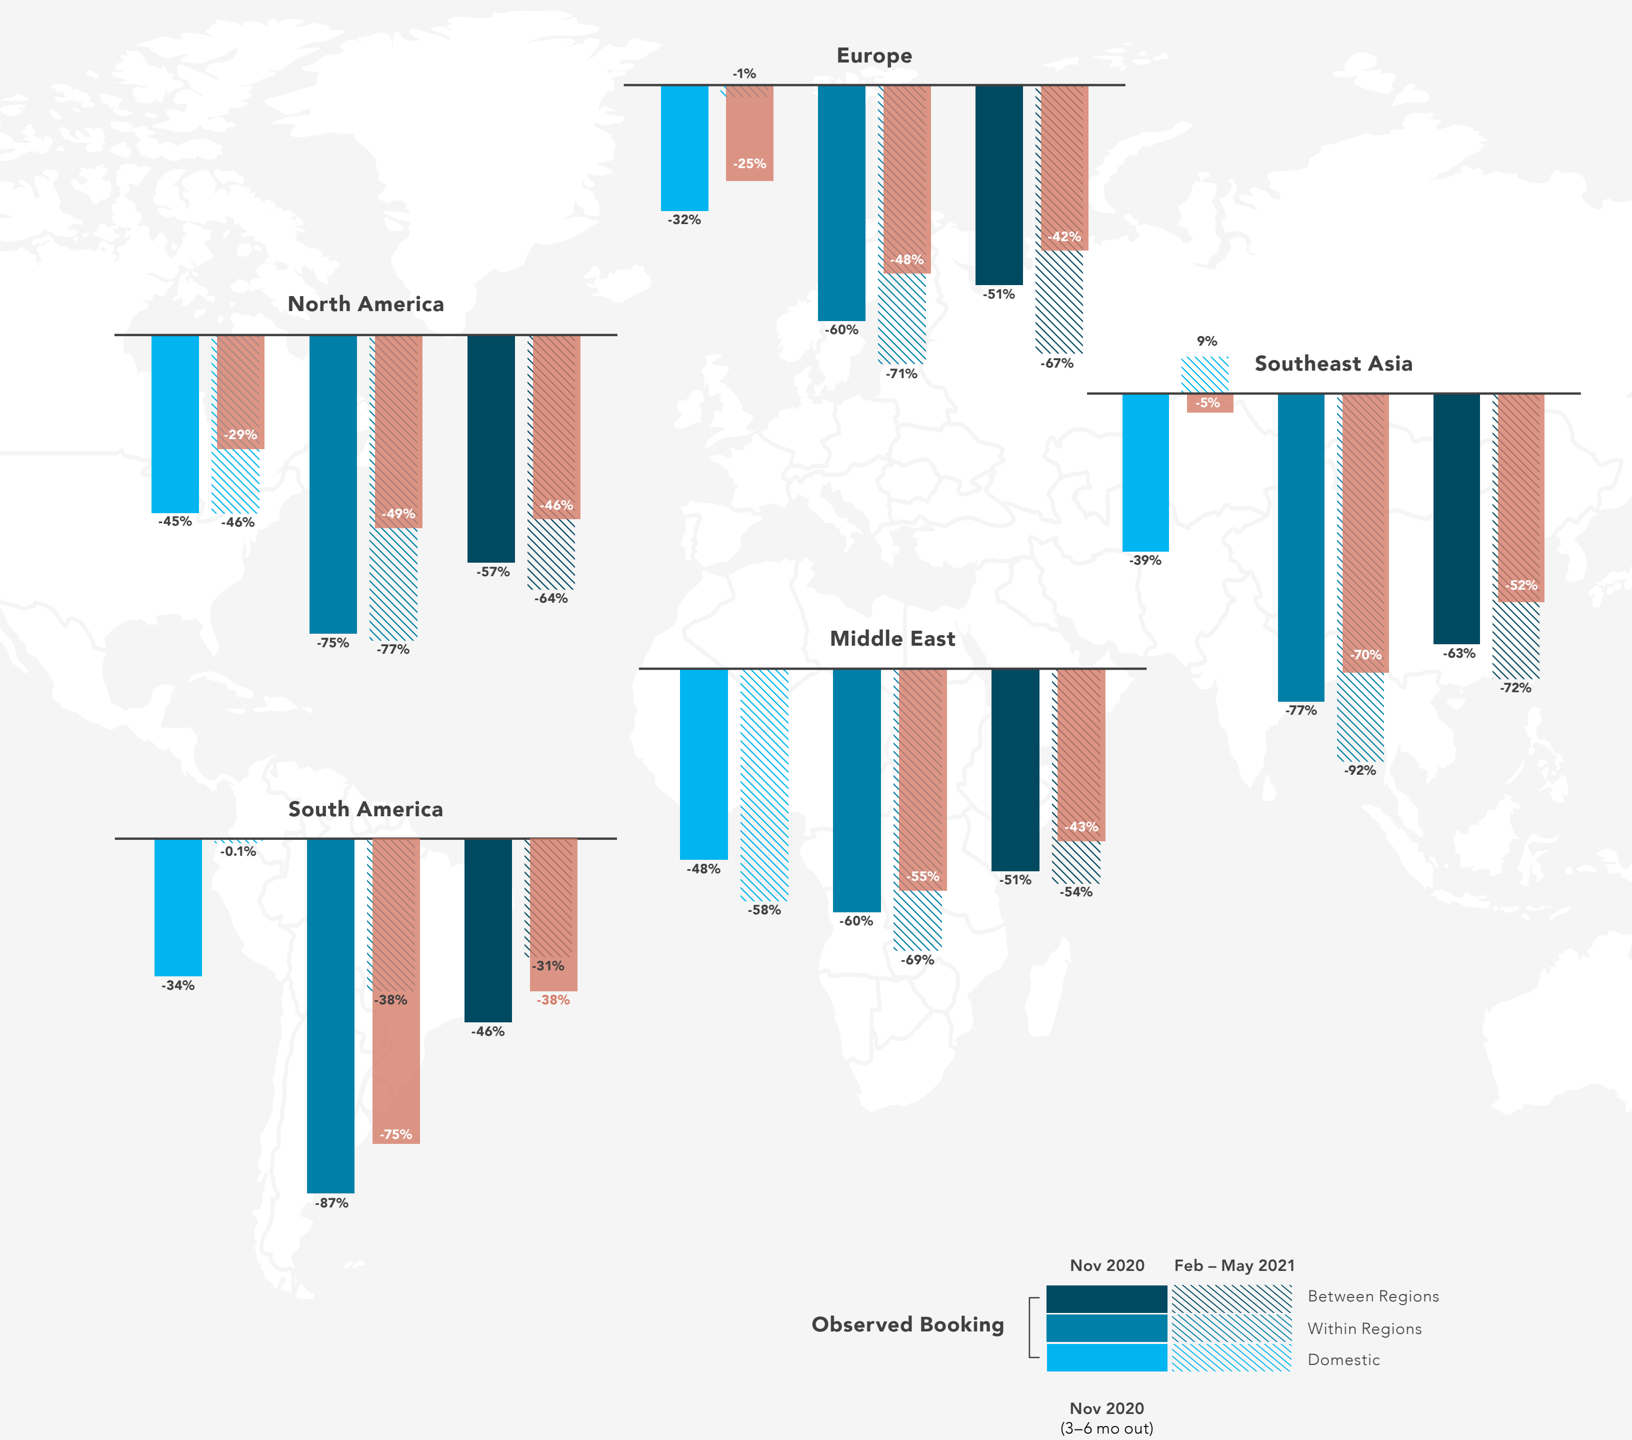


FIGURE 3

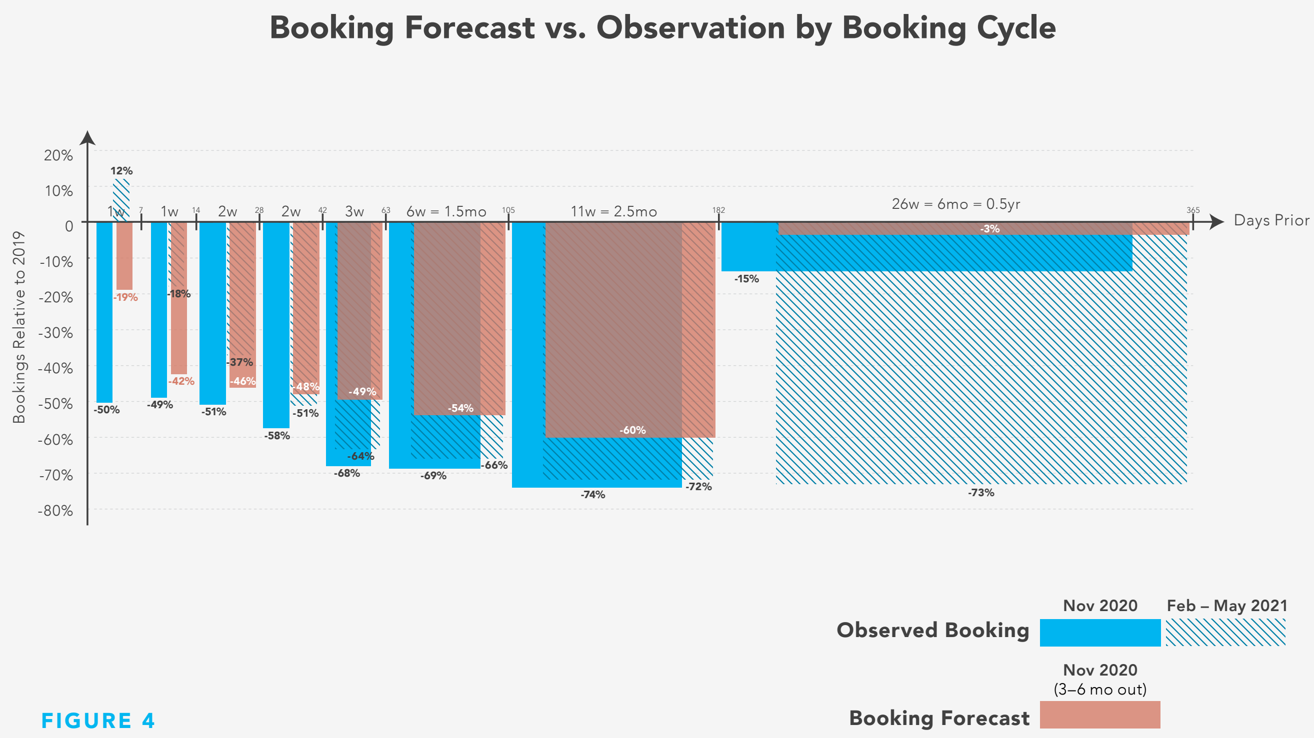


**FIGURE 4**

At first glance, we clearly missed the mark! But considering the limited alternative at hand, this wasn't terrible. One of the reasons for this suboptimal result is that we couldn't have anticipated the appearance of new variants (see Table 4 below) with increased infection rates.

| COVID Variants | [Delta]() | [Gamma]() | [Omicrom]() |
|---|---|---|---|
| Date of earliest sample | Oct 2020 | Nov 2020 | Nov 2021 |
| Date designated as variant of concern | May 6, 2021 | Jan 15, 2021 | Nov 26, 2021 |

**TABLE 4**

These more infectious variants will have a differential impact on both government response and consumer sentiment on travel safety. Even IHME's forecast of the pandemic's progression were off by a similarly wide margin. Besides, our booking forecast is an extremely challenging problem, because we are trying to forecast far out into the future for a 3-month long window (3 to 6 months from Nov. 2020).

Consider the weather forecasting problem. With over a thousand-year of meteorological science and innovation, weather forecasting is a very mature discipline. And most people would consider today's weather forecasts to be very accurate. With ubiquitous sensor networks, we can even forecast accurately down to the minutes when the rain will start or stop. However, if we were to forecast the weather far in the future (e.g. 3 months from today) for a lengthy window (e.g. 3 months long), the performance probably wouldn't be much better than ours. Despite the inadequacy of weather forecast, it is still extremely valuable.

Another reason that forecasting problems are so much harder than prediction problems is that forecasting is inherently time-dependent. For example, the weather forecast for New Year's Eve (NYE) is a quantity that will change as we move closer and closer to NYE. Because as we march forward in time towards NYE, not only will we have more data with each passing day, we also have more accurate data. Moreover, the shorter time-horizon to the forecast window means fewer things can change in the meantime. So the weather forecast for NYE 2-weeks prior may be very different from the weather forecast for that same day just 1-week before.

Therefore, most forecasting models are run repeatedly at some pre-determined frequency, and the forecast in the target window will improve continuously. Moreover, it is crucial to continuously refine our model since the underlying forces that drive demand are constantly and rapidly changing. Since our last forecast, vaccines were deployed and became more accessible, yet new COVID variants also appeared. Not only will different countries respond differently, since the mortality rates of these variants are different, their impact on consumers' travel demand also varies greatly.

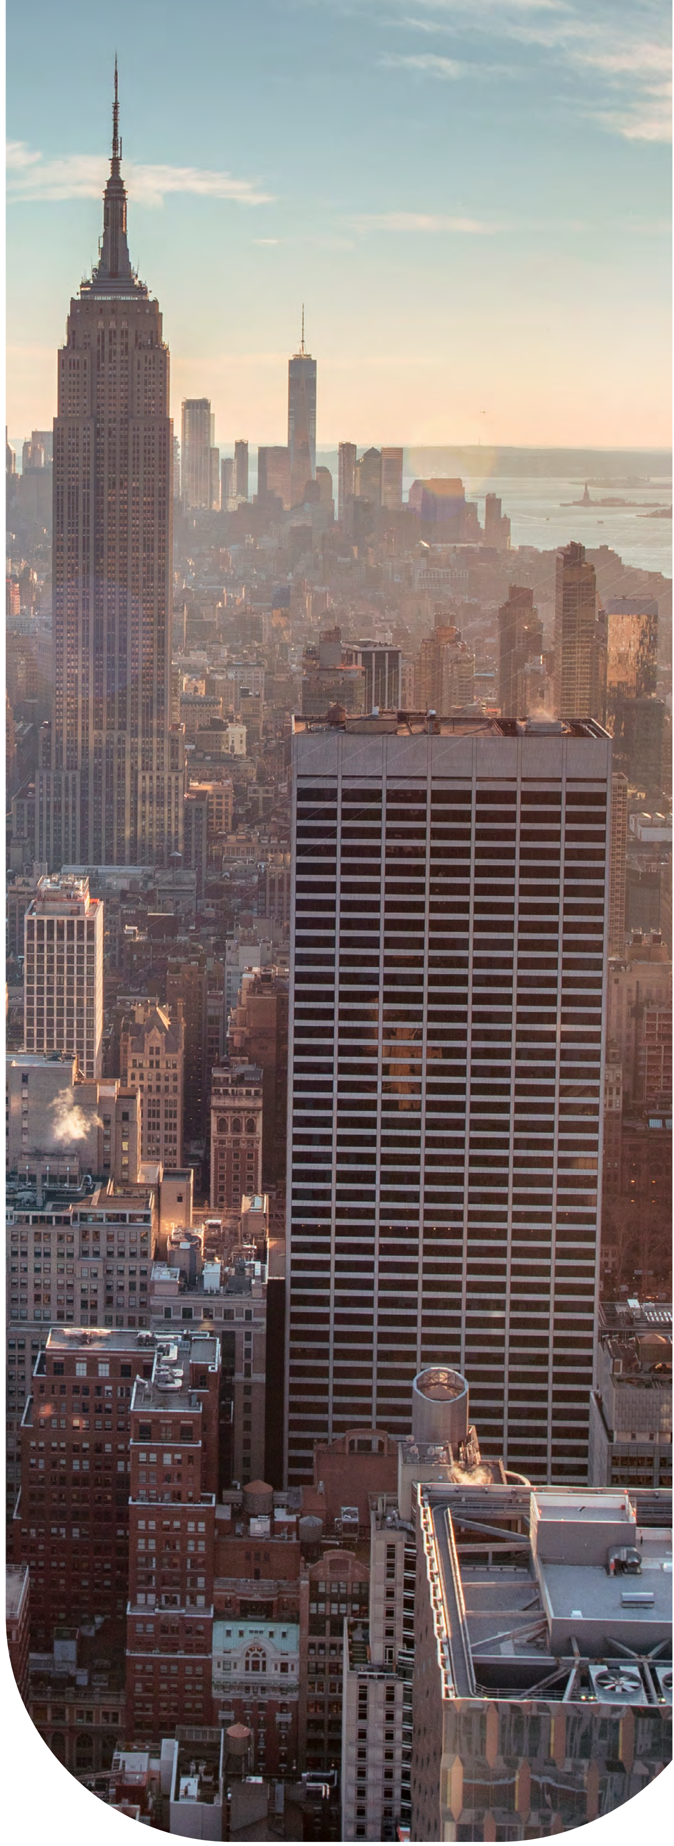

# Our Current Booking Forecast as of Oct. 2021

We've strived to continuously refine our model to keep up with the pandemic as it evolves. Since our last booking recovery forecast in Nov. 2020, we have greatly improved our model. We've leveraged new data sources (e.g. vaccination data) as well as engineered new features. Using the Pearson's correlation coefficient (PCC) as a prediction accuracy metric, the performance of our latest random forest model is able to predict with an average PCC of 0.68 with the best PCC at 0.74. Since PCC is normalized between –1 and 1, our refined model is performing very well.

So, what could this enhanced model tell us about the future demand for travel? In Oct. 2021, we produced another booking recovery forecast. Due to our agreement with the CTF participants, we are only allowed to publish these forecasts at the regional level. In the following discussions, all forecasts are normalized to 2019 level to obfuscate the market shares of our airline participants.

The first set of forecasts is the passenger mix of the returned bookings (see Figure 5). As of Oct. 2021, booking across all cabins are about 25% – 35% below 2019. Our model suggests a modest improvement of ~7% for both the first class and economy class cabin in 3 – 6 months. However, as the global workforce becomes more accustomed to remote work and more business gatherings shift to virtual, business class bookings are forecasted to fall ~6% further below 2019 level.

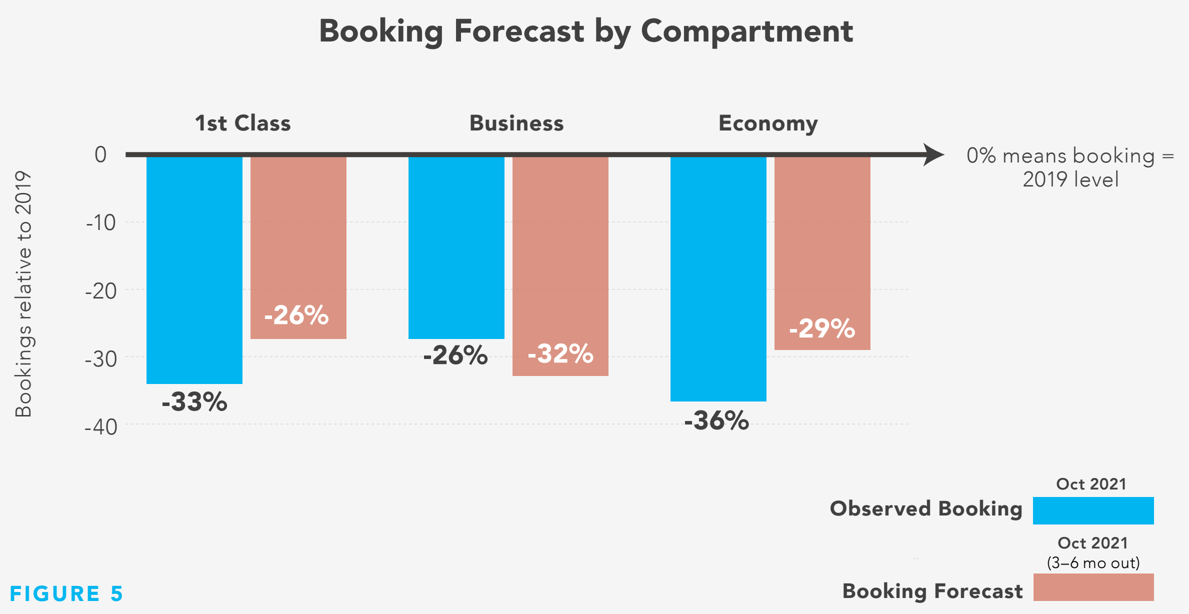


**FIGURE 5**

The next set of forecasts is based on where the passengers are and where they are going. We've defined seven regions for our study: North America, South America, Europe, Africa, the Middle East, Southwest Pacific, and Southeast Asia. However, we did not produce a forecast for Southwest Pacific and Africa, because the bookings in these two regions were below the threshold needed for a reasonably accurate forecast. For the remaining five regions, we have analyzed three types of flights:

1. **Domestic flights**
2. **International flights within a region (intraregional)**
3. **International flights between regions (interregional)**

We will highlight some interesting observations among all the data in Figure 6 on the next page. As in the previous forecast, domestic flights have been recovering the most among the three flight types, except in the Middle East. In fact, domestic flights in Europe have risen above 2019 level by ~8% as observed in Oct. 2021. In 3 – 6 months, North America domestic bookings would improve significantly (by ~23%), whereas Southeast Asia would only improve slightly (by ~3%). On the contrary, domestic bookings in Europe are forecasted to fall ~20% below the current level, to ~13% below 2019 level. This is presumably due to the emergence of the highly infectious Omicron variant that is spreading rapidly across Europe.

We did not provide a forecast for domestic flights in South America or the Middle East, because our data do not have a sufficiently diverse pool of carriers that offer domestic flights in these regions. So, our forecast is not generalizable for all the countries within these regions.

## Booking Forecast by Region and Flight Type

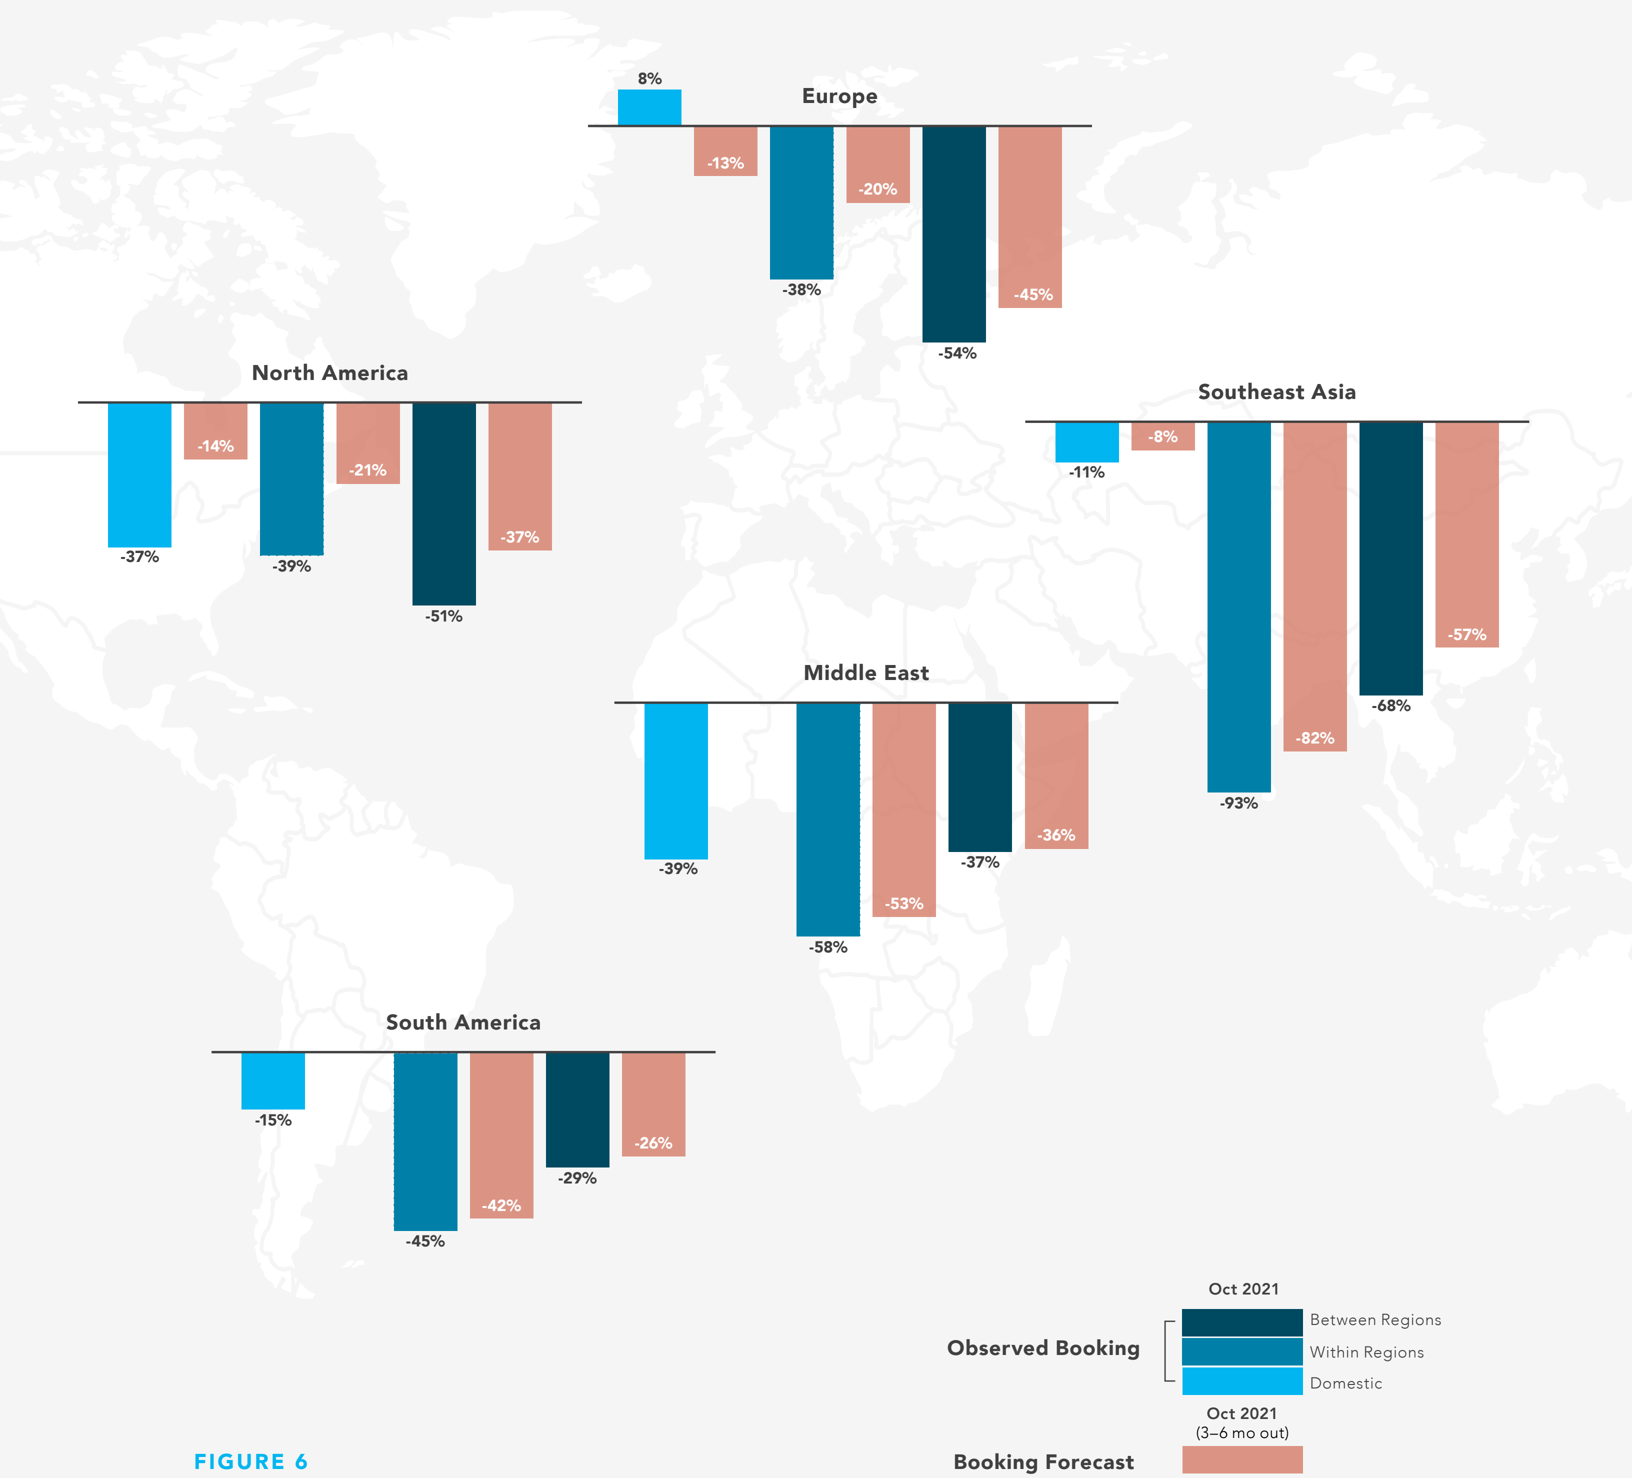


FIGURE 6

North America and Europe are seeing a similar recovery pattern with their intraregional bookings (~40% below 2019 level) recovering ~10% more than their interregional bookings (~50% below 2019 level). And in 3 – 6 months, the intraregional bookings are forecasted to further improve by ~20% to just ~20% below 2019. But their interregional flight is forecasted to have a more modest improvement of ~10% to a level that is ~40% below 2019.

In general, international bookings in South America, the Middle East, and Southeast Asia are seeing a slower recovery in their bookings compared to North America and Europe as of Oct. 2021. But unlike North America and Europe, their interregional bookings are recovering more than their intraregional bookings. However, their booking forecast in 3 – 6 months would only improve slightly (~10% or less) across these three regions. International flights in Southeast Asia are forecasted to improve the most (~10%) even though they are currently seeing the slowest recovery compared to 2019. International bookings in South America and the Middle East are forecasted to improve less than 5% while they are currently recovering more than Southeast Asia.

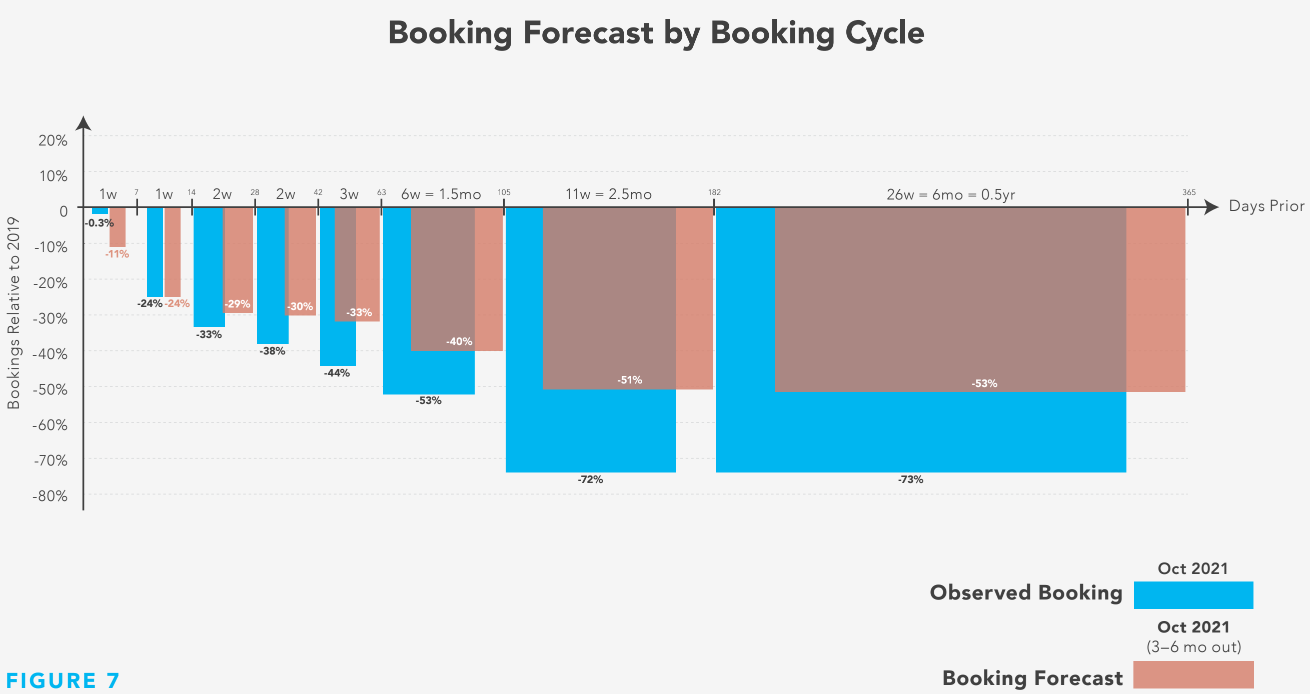


**FIGURE 7**

The last set of forecasts is about people's propensity to travel in the future. Since people can book flights up to one year before their departure date, and there tend to be much fewer bookings for a further-out travel date, it's a common practice to bucket the number of bookings in days-prior to the travel date into timeframes of different length. These timeframes are represented as different bin sizes (i.e. 1 week, 1 week, 2 weeks, 2 weeks, 3 weeks, 6 weeks, 11 weeks, 6 months) in Figure 7.

As of Oct. 2021, bookings for departures within 1 week have completely recovered to 2019 level. But it is forecasted this will be worsened by ~10% in 3 – 6 months, so the booking level will be ~11% below 2019.

Bookings for departures that are more than 1 week out are still far below 2019 level. As the travel dates move further out, the observed booking as of Oct. 2021 falls steadily from ~25% below 2019 and reaches a level that is ~70% below 2019 for travels beyond 3.5 months (105 days) in the future. Despite the grim outlook with the current situation amidst rising COVID cases, the good news is that this will improve substantially in 3 – 6 months. The forecasted bookings (in 3 – 6 months) will improve progressively more with further-out departures, and bookings that are more than 3.5 months out could improve as much as ~20%.

# Optimizing Your RM for the Post-Pandemic New Normal

With an improved model and improved forecast (prediction about the future), this is where most of the academic research ends. But this is usually not enough for businesses, because even when they know how the future will unfold, there are still virtually an infinite number of ways they can respond to it. Should you respond to this forecasted future in the fastest way, the most effective way, the cheapest way, or just any way that beats your competitor? Or perhaps the best course of action is simply to do nothing at all. So, forecasts are not sufficient for businesses because they may not be actionable. What businesses need is often not a prediction of the future, but rather a prescription of precisely what to do about it.

To make our booking forecast more actionable, we must translate this knowledge of the future into a specific set of actions. Fortunately, this can be done by tuning our RM systems to function optimally under the forecasted conditions. During disruptions, one of the most important variables we must optimize for any AI solution that learns is its learning rate. Learning too slow will obviously hinder the system from adapting quickly, but learning too fast will make the system unstable.

For airlines' RM, this learning rate is controlled by the drift parameters of the Bayesian forecaster, which control the exponential time decay for the impact of historical data. So, a good set of drifts should allow the forecaster to react just fast enough to track the recent booking fluctuations. We've implemented two methods to determine the optimal drift parameters. Detailed exposition about these methods can be found in our CTF blog update #6. But, in short, the first method is a brute force minimization of the forecast error as a function of the drift parameters over the recent histories. And the second method is based on reverse engineering heuristics.

Since the drift parameters in the PROS RM system are global settings, they set the overall learning rate for our Bayesian forecaster. But how can this system be optimized for different markets

with different levels of disruption and different recovery patterns? This can be accomplished indirectly by censoring, as we can apply data censorship at a very granular level, down to the origin and destination (O&D) pairs and below. Determining the best censoring strategy is nontrivial, because it's a delicate balancing act and depends on what each carrier has already implemented for each market. Again, our method is discussed in much greater depth in our CTF blog update #6.

The output deliverable of this study is nothing flashy. The CTF participants will get a few numbers for the updated drift parameters and a list of censoring recommendations for their O&Ds. However, the true benefit of being a part of the CTF is the ability to adapt to the new normal faster and recover from the pandemic sooner. And this is how we address the category 3 inquiries on how to reset for the new normal.

# The Objective Evidences of a Positive-Sum Game

So how have we done? Is our work helping our customers and adding value to the CTF participants? Can we answer the question we posed initially on "whether a large-scale collaboration among the carrier is creating greater value to everyone?" After 20+ long months, the CTF has two major outputs as deliverables back to the airlines.

1. **The forecasted global demand recovery pattern**
2. **The recommended tuning parameters for each airline participant**

We can objectively quantify the performance of these outputs against real data. Furthermore, we can subjectively quantify the value of CTF by surveying the roughly two dozen airline participants in the study. The first objective performance metric is the percent-error reduction of the joint model (trained with the combination of all the airline participants booking data) relative to the single-carrier model (trained with only the airline's own booking data). Essentially, we want to know how much better the joint model can forecast future demand compared to the single-carrier model (see Figure 8).

**Performance Metric 1**

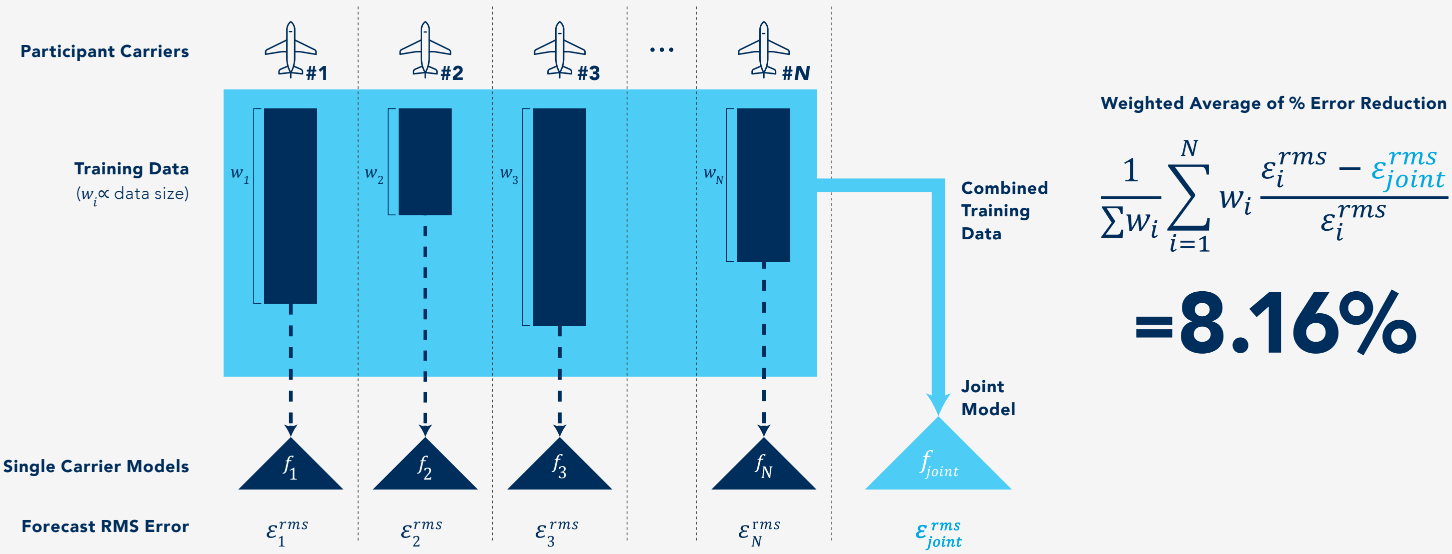


**FIGURE 8**

To mitigate the effect of accidental temporal correlation, we have chosen two different test periods to validate our model. Moreover, we have computed the error reduction two different ways and obtained the expected error reduction on average. Evidently, the joint model's forecast has a root mean square error (RMSE) that is on average 8.16% lower than the forecast from the single-carrier model. So, the joint model can forecast 8.16% better than the single-carrier model on average as seen by each airline.

The second objective performance metric is harder to obtain, since it requires the participating carriers to have an existing censoring strategy that is compatible with our recommendation. Essentially, we are implementing our recommendation in a non-production environment and performing an experiment *in silico*. We can then quantify the performance gain of our recommendation by computing the forecast error reduction under the new settings relative to the forecast error under existing settings (see Figure 9).

## Performance Metric 2

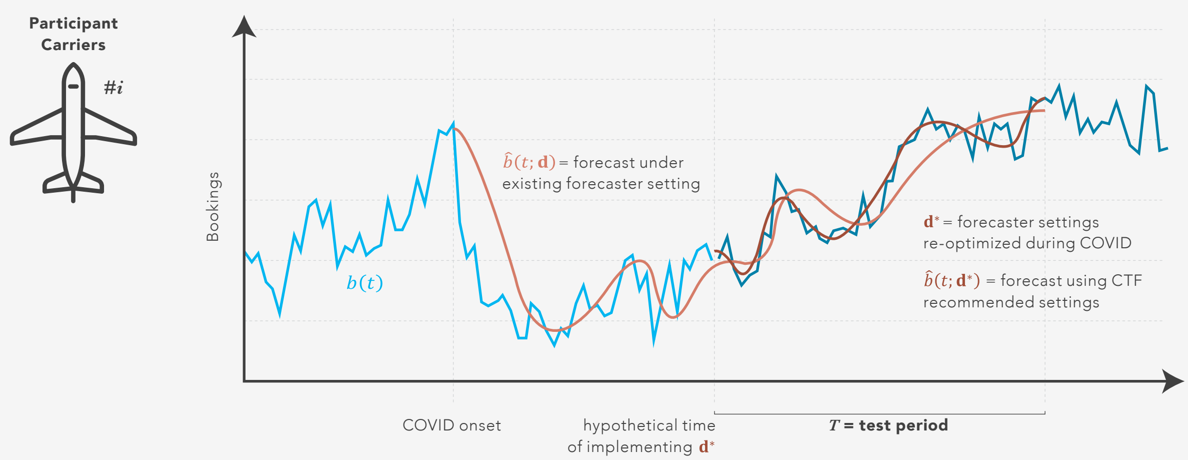


RMS forecast error under existing forecaster setting

$$\varepsilon_{i-}^{rms} = \sqrt{\frac{\left[b(t) - \hat{b}(t;\mathbf{d})\right]^2}{T}}$$

RMS forecast error using CTF recommended forecaster setting

$$\varepsilon_{i+}^{rms} = \sqrt{\frac{\left[b(t) - \hat{b}(t;\mathbf{d}^*)\right]^2}{T}}$$

**Average % RMS Error Reduction**

$$\frac{1}{N}\sum_{i=1}^{N} \frac{\varepsilon_{i-}^{rms} - \varepsilon_{i+}^{rms}}{\varepsilon_{i-}^{rms}} = 6.37\%$$

**FIGURE 9**

Furthermore, we must quantify this forecast error reduction for a sufficiently long period after implementing our recommended settings to ensure the performance gain is not transient. At the time of writing of this whitepaper, we can only compute this performance metric for 3 carriers with existing censoring strategies that are compatible with our recommendation. And the average expected RMSE reduction after implementing our recommendation across these 3 airlines is 6.37%.

# The Subjective Values of the COVID Task Force

Beyond the definitive positive impact of combining data across the airlines, we also surveyed the participants for their opinion about the value of the CTF's work. In general, we are glad that the participants are satisfied with our work, as we receive **3.83 out of 5 stars in terms of overall customer satisfaction.**

However, since value is always in the eye of the beholder, different carriers value different aspects of our work. The survey data suggest that most airlines found our webinars and the optimized drift and censoring recommendation most valuable (see Figure 10).

**In fact, 83% of the survey respondents said they have either implemented our recommendations already or plan to implement them soon.**

### Relative Utilities of COVID Task Force Deliverables

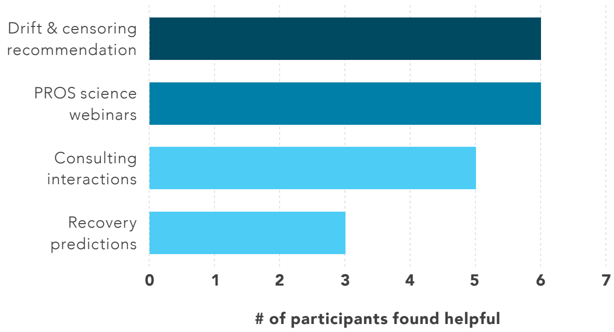


FIGURE 10

Despite the fact that CTF participants are anonymous, some participants were willing to provide us with a referenceable quote to validate our effort. Here's what some of our airline participants are saying about our work:

*"This pandemic brought revenue management systems and practice to an abrupt halt at most airlines. Leveraging a unique position in the airline industry, PROS was able, through the COVID Task Force, to explore integrating non-conventional signals and metrics, and testing alternative calibrations and schemes aiming to 'jump-start' the passenger demand forecast engine.*

*At Air Canada, the Task Force provided insights which helped us craft a system-deconfinement strategy to strike the right balance between censoring and allowing the system to learn again and provide a solution adapted to the new normal."*

**Richard Cléaz-Savoyen,**
Senior Director, Revenue Optimization
AIR CANADA

*"As a participating carrier in the COVID-19 Task Force Study, Emirates Airlines was able to leverage the data science expertise of the PROS CTF to enable the PROS RM system to adapt faster to the new normal demand patterns.*

*The PROS team performed a customised forecast accuracy study illustrating the expected impact of the optimised drift and censoring recommendations enabling the implementation of the required changes."*

Emirates

# Demand Forecasting Beyond Airlines

In addition to the value we created for the participants, the CTF has provided an opportunity to co-innovate with our participants. Although the CTF participants were primarily airlines, we did have a small number of non-airline participants spanning the transport, hospitality, and food industries.

Due to the limited number of industry-specific participants, we were unable to replicate our large-scale airline study with non-carrier participants and obtain statistically significant conclusions. Therefore, the depth and nature of the investigation with our non-airline participants were individualized, even though the goal is still aimed at understanding demand recovery from the pandemic.

Because some of these non-airline participants have provided us with additional data sources, we were able to launch two spin-off research initiatives.

1 **We have investigated the use of shopping data to improve the accuracy of demand forecasting.**

2 **We also experimented with the use of events data in estimating the impact of events on consumers' shopping behavior.**

A detailed exposition of either of these initiatives is beyond the scope of this whitepaper. Although these researches are still in their early stages, the preliminary results are very promising.

**Using data from a transport and logistics company, we found shopping data is able to predict demand with an average accuracy of 0.54, but can go as high as 0.97 as measured by PCC.**

The booking horizon for this transport and logistics company is much shorter (only about 2 weeks) than airlines, however. That means people typically do not shop for or book until just few weeks before the service is needed. So even though shopping data can accurately predict the demand for transport and

logistics services, we cannot use it to forecast demand far into the future, because such far in advance shopping data does not exist.

This is why we started the second research initiative to investigate the use of event data to predict shopping behaviors. Because events are typically planned and announced far in advance, an accurate prediction of shopping from events implies that we can get a good proxy for people's shopping behaviors far in advance. This estimated shopping activity can then be used to forecast booking far into the future, way before people even start their shopping (i.e. before shopping data materializes).

Using events data from a third-party provider and shopping data provided by our CTF participants, our research demonstrates that:

**We can predict shopping activities with an average correlation of 0.59. But in the best-performing markets, the PCC can go as high as 0.72.**

Although these research initiatives began with our non-airline participants in the CTF, they do have important use cases for airlines. We are currently expanding the scope of these investigations to see if we can replicate these positive results for airlines. Initial results suggest that airline shopping data, such as flight search data, can accurately predict future booking trends. Therefore, these research initiatives could greatly improve our demand forecasting engine and potentially replace the user-defined holiday & special events module.

Furthermore, the use of events data could even help airlines allocate their capacity earlier in the network planning stage. We have briefly outlined a tentative plan on how to leverage these research findings to innovate the future of RM in our final CTF blog.

# Protecting Our Partners From Beginning to End

Like the pandemic, the CTF is also an effort that is unprecedented in the history of PROS. It is the first time we have had so many customers sharing their demand data to be combined with those from other CTF participants, some of which may be direct competitors. So, we are very diligent in protecting the trust and relationships with our customers. We had to create new processes and analytics infrastructures to handle customer demand data. Even our legal and governance team had to get involved to amend our contract with customers on how their demand data will be used, analyzed, and purged after our work is complete.

This demand data is extremely sensitive as it includes bookings for the airlines, and it may contain volume, frequency of purchase, or, in some cases, revenue for our B2B customers. We have formulated all our problems based on correlation and regression, so the absolute number of bookings or the exact dollar amount of revenue is not important. This will allow us to perform the analyses on a rescaled version of the highly-sensitive demand data.

Whenever bookings from one airline need to be intermingled with bookings from another airline, they are pre-aggregated to the country level before they are combined and then analyzed globally. To further protect the identity of our participants, all demand-related predictions will be reported in relative terms. All results will be expressed in percentages or percent changes relative to another part of our customer's own data to prevent any identifiable information from being reverse engineered.

As we are close to the end of 2021, we are wrapping up with all the promised deliverables to our CTF participants. All our investigations are complete, culminating in the production of this white paper. So, it is time to adjourn the CTF and erase all customer demand data. Over the past 20+ months of CTF investigation, we have accumulated over 2 TB of booking data from about two dozen airline participants. Keep in mind that 2 TB is a huge amount of data since it consists of mostly numbers and text, not images or videos. If we print this data out, it would make a stack ~10km high and weight ~500 metric tons. This data constitutes approximately 731 million booking records covering 129 countries.

## With a single keystroke, in a fraction of a second, poof! All that data... gone!

Suddenly, I felt a certain inexplicable ache. How can we so easily destroy so much data that we diligently collected over the years? Even though this data is technically not of our own, we treated it like our own and protected it like our own. Purging data has always been stressful and very uncomfortable for someone with data-separation anxiety. The little and only comfort I drew from this sudden loss of data is the fact that we've gained valuable knowledge from it.

# Concluding Remarks

Although the air travel industry is extremely competitive, it is comforting to see that in dire times, we can set aside our difference and collaborate. This not only requires a leap of faith in the airlines, but also a huge investment from PROS.

We have invested approximately 8,000 person-hours, which translates to a full-time data scientist working for about 4 years. This doesn't even include the time investment from departments like legal, data governance, cloud IT, support, marketing, etc., not to mention the infrastructure cost required to store, secure, manage, and process all the data.

The insight from all this is compelling evidence that even in the highly competitive airlines industry, productive collaboration through data sharing can happen. And it is possible to turn their zero-sum games into a positive-sum game.

**Our analysis suggests that when airlines combined their booking data, they can forecast future demand 8.16% better on average as measured by RMSE reduction.**

**In addition, implementing our re-optimized RM configurations can improve demand forecasting accuracy by 6.37% on average.**

So, we can collaborate to grow a bigger pie for everyone, and we can objectively quantify the benefit of this collaboration.

Now we have the evidence that even extremely competitive industries, like the airlines, are not zero-sum games. In fact, the positive network effect will always make the sum greater than its parts. What's left now is for the airlines to take the next step and find new and innovative ways to collaborate. We would be honored to facilitate this collaboration. I just hope that it will not take another global disruption, like the pandemic, to instigate this. Because I strongly believe that if we choose solidarity, we will all win. And together, we will all return stronger to compete again in the next big race.

# About the Contributors

This work is the culmination of many scientists in the CTF with a substantial contribution from many different departments across PROS.

This is truly a heroic effort from everyone, and we would be superheroes in the metaverse. Although it would be impractical to list everyone who has contributed to the CTF, here are our brilliant scientists alphabetically ordered by their last name. We also included a quote from each one of them about their unique contribution to the CTF.

**AMY BOCK** "I did a little bit of everything on the CTF and nothing without my teammates' help. This includes data operations, customer study deliveries, conceptual modeling, and turning our predictions into forecasts."

**SHAHIN BOLUKI** "I worked on enhancing the prediction and forecast models. I also worked on the prescriptive analytics and designed methods for tuning the RM system."

**LALEH KARDAR** "I contributed to CTF by performing the forecast accuracy study to show how much forecast accuracy would improve if airlines implement task force recommendations."

**JIABING LI** "I work on building and refining a general shopping model, which leverages shopping data to improve booking forecasts."

**YARA LIU** "I work on the machine learning pipeline. Specifically, I tune and optimize the hyperparameters of our random forest model."

**MALORI MEYER** "I contributed to the project throughout the full pipeline with the bulk of my work building the ETL pipeline, training and tuning the random forest model, and the forecasting process."

**FARUK SENGUL"**My contribution to CTF is data preprocessing (cleaning, mapping, observation building, and aggregating) and post processing. I also helped review the shopping model."

**WENSHEN SONG** "I mainly contribute to the CTF by developing a model that estimating the impact of events on consumers' shopping behavior."

**ROSS WINEGAR** "I was the Team Lead for the CTF. Bridging between our participants and our scientists, I was involved in all aspects of the CTF, guiding the team from the beginning up through the deliverables."

**MICHAEL WU** "I am the chronicler of the CTF, responsible for communications and marketing. I also contribute to model development, data visualization, and presentation."

**YAN XU** "I contribute to the CTF by providing advice on how to configure our B2B pricing suite to better adapt to the market disruptions due to COVID."

**ZHAOYANG ZHANG** "I helped my teammates tuning the drift parameters for our clients."

PROS Holdings, Inc. (NYSE: PRO) is a market-leading provider of SaaS solutions that optimize shopping and selling experiences. Built on the PROS Platform, these intelligent solutions leverage business AI, intuitive user experiences and process automation to deliver frictionless, personalized purchasing experiences designed to meet the real-time demands of today's B2B and B2C omnichannel shoppers, regardless of industry.

**To learn more, visit pros.com.**