KRANNERT SCHOOL OF MANAGEMENT

Highlighting Management Department Faculty Research

CONTENTS :

Management Department Head
Ananth Iyer

We are pleased to present the 4th research newsletter, showcasing management faculty research across all eight academic areas, published by the management department at the Krannert School of Management at Purdue University. If there is one theme in all of these papers, it is the dedication of our faculty to scholarship, to their drive to understand, synthesize and recommend actions in their domain of interest. I hope you learn from their insights, and provide us feedback regarding this newsletter. We will continue to strive to communicate our faculty's research to our stakeholders.
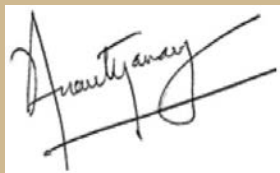
This newsletter contains a diverse set of topics. Diana Choi explores if including finance experts in bank audit committees "improves the timeliness of loan loss provisions, thus ensuring that they quickly reflect changes in credit risks". Ha Diep-Nguyen studies how "creditor-shareholder conflict affects the resolution of financial distress", a topic that is particularly relevant as firms assess the impact of the pandemic on their finances. Will Haskell's research proposes "methods that are sufficiently universal that they can produce good strategies using only data readily available about the problem". Yixuan Li's research helps uncover the mechanisms through which age diversity affects organizational performance.

Luis Rios's research shows that firms that persistently choose alliances and joint ventures outperform those favoring acquisitions and divestitures vehicles during stable economic periods, but the reverse holds true in recessionary periods.

Yixuan Liu studies competition and participation in on-demand healthcare platforms, which have been growing in popularity, particularly during the coronavirus pandemic. Yichen Zhang's research develops a fully online approach to construct asymptotically exact confidence intervals for large-scale data sets. He cites an example of a pharmaceutical company developing a new drug that needs to make a hypothesis test to determine whether the drug is valid to treat a specific illness. Guofang Huang's research proposes an approach to analyze survey data that adjusts for bias, and thus suggests that many companies may have "grossly underestimated" the return on investment in customer service.

This research newsletter provides you but a glimpse of the intense focus on research excellence and the clarity of the problem framing by our faculty, as well as the care with which results are presented to ensure timeless value. We hope that you are motivated to look up the associated academic papers, or connect with the faculty for more detail. But, more importantly, please do not hesitate to seek out the expertise of our faculty to provide you with insights regarding your own business environment. We hope that this newsletter provided you with new ideas to ponder, and that you to continue to read our research newsletters and provide feedback.

Have a wonderful day.

Ananth V Iyer, Department Head,
Susan Bulkeley Butler Chair in Operations Management
Krannert School of Management
Purdue University
West Lafayette, IN 47906

Financial Expertise is Important on Bank Audit Committees, But so is Independence

Diana Choi

Simultaneous Debt-Equity Holdings Reduce Conflict, Help Firms Avoid Bankruptcy

Ha Diep-Nguyen

When Real-World Problems Have Too Many Potential Outcomes, Universal Randomized-Function Algorithms May Help Decision Makers

William Haskell

How Age Diversity Can Help Organizations Perform Better

Yixuan Li
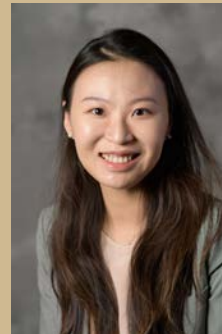
Doctors Competing Against Each Other: What's Good for Patients May Not Be Good for On-Demand Healthcare Platforms

Yixuan Liu

A Better Way to Measure the Impact of Customer Service Satisfaction on Loyalty

Guofang Huang

Acquisitions or Alliances: Why Companies Persistently Choose One Path over Another

Luis Rios

Statistical Inference for Stochastic Gradient Descent: A Fully Online Approach to Construct Confidence Intervals

Yichen Zhang

Financial Expertise is Important on Bank Audit Committees, But so is Independence

by Melvin Durai

Banks make loans to a variety of customers, ranging from individuals with few assets to large corporations worth billions. The borrowers' financial status impacts loan quality and determines whether banks are assuming low, moderate or high risks. To account for potential loan defaults and expenses, banks are required to include loan loss provisions in their financial statements, thus providing a more accurate representation of their financial health.

While bank managers determine provisions for loan losses, it falls on audit committees to monitor this process and verify that these highly subjective estimates are reasonable. Most banks appoint financial experts to their audit committees, complying with rules designed to reduce conflicts of interest. But do these experts improve the timeliness of loan loss provisions, ensuring that they quickly reflect changes in credit risks?

That's a question that Diana Choi, assistant professor of accounting in Purdue University's Krannert School of Management, has helped answer through her research.

She found that while financial experts on audit committees improve the timeliness of loan loss provisions, the impact varies according to their background. Most notably, former regulators and affiliated auditors are associated with less timely loan loss provisions.

"I was surprised," Choi said. "I actually hoped to find that those external auditors or former regulators would bring more expertise to the audit committees of banks, so they would promote more timely provisioning, but I find the opposite result, which indicates that they have different incentives."

Her findings, shared in a paper entitled "The Effect of Bank Audit Committee Financial Experts on Loan Loss Provision Timeliness," suggest that banks should consider more than just financial expertise when making appointments to audit committees.

"If you agree that financial reports are important to investors or anyone who is a decision maker, having an audit committee that has certain financial expertise is very important," Choi said. "But one thing we should not neglect is whether the financial experts are truly independent to make sure that the financial reports provide timely information."

To examine the impact of financial experts sitting on audit committees, Choi used the Federal Depository Insurance Corporation Improvement Act of 1991 (FDICIA) and a listing standard for NYSE and NASDAQ firms in 1999 as quasi-natural experiments. The FDICIA, passed during the savings and loan crisis, required that the audit committees of banks with total assets greater than $3 billion include members with relevant banking or financial experience. The listing standard required that firms listed on the two exchanges have at least one director with accounting expertise on their audit committees.

Both of these rules were created to address agency problems between managers and stakeholders, giving financial experts the role of mitigating these conflicts of interests. These experts might be expected to help bank management identify potential losses quickly and repress any reluctance to recognize them in financial statements in a timely manner. But the converse could happen if these experts strictly follow accounting models that recognize losses only when they are probable.

To determine what was actually happening, Choi employed a difference-in-differences research design, comparing two groups of banks. The first group consisted of banks with assets greater than $3 billion and thus falling under the FDICIA rule requiring financial experts to sit on audit committees (the treatment group), and the second consisted of banks with assets between $500 million and $3 billion (the control group). Each group included a sample of banks traded on the stock exchanges, allowing Choi to test the NYSE and NASDAQ listing standard.

Focusing on the three years before and after the regulations requiring financial expertise on audit committees were instituted, she found that the rules increased the timeliness of loan loss provisioning. While her FDICIA treatment group consisted of larger banks than the control group, her NYSE and NASDAQ treatment group consisted of smaller banks than the control group, alleviating concerns that the size of the banks influenced the results.

To examine the effects of appointing different types of financial experts to bank audit committees, Choi considered former affiliated auditors, former unaffiliated auditors, former regulators, and those who were neither auditors nor regulators.

"'Affiliated' indicates that affiliation exists between banks and external audit firms that bank auditors worked for," she said. "What I did was identifying those audit committee members who used to be external auditors in the audit firms and then I identified whether the banks were audited by the same audit firms."

She found that former affiliated auditors are associated with less timely provisioning compared to unaffiliated auditors, while former regulators are associated with less timely provisioning compared to unaffiliated auditors and those who were neither auditors nor regulators.

"It is hard to find those people who have good financial expertise, but are also truly independent," she said. "It is difficult, but it is something that banks should aim for to make sure they could have financial reports that reflect their loan quality."

Simultaneous Debt Equity Holdings Reduce Conflict, Help Firms Avoid Bankruptcy

By Melvin Durai

The coronavirus pandemic has pushed the economy into recession and shrunk the revenues of many firms. Some will be forced to file for bankruptcy, while others may be able to achieve a successful out-of-court restructuring or workout, reaching an agreement with significant creditors to adjust obligations.

While a workout is often more attractive, helping a firm lower costs and retain more of its value, it's not always possible, partly because creditors may be unwilling to participate in an out-of-court restructuring. Compounding this are the diverging interests of creditors and shareholders.

"Creditors, especially those that are protected by collateral, don't get hurt if the firm goes into bankruptcy," said Ha Diep-Nguyen, assistant professor in Krannert School of Management. "On the other hand, if the firm continues, instead of going into bankruptcy, the creditors may be concerned that it will take on too much risk. For them, bankruptcy may be a safer bet. But for shareholders, their interests are often wiped out in bankruptcy. If the firm can restructure out-of-court and regain its value, the shareholders can enjoy the benefits of this."

Research by Diep-Nguyen and her collaborators, Yongqiang Chu of University of North Carolina at Charlotte, Jun Wang of University of Western Ontario, Wei Wang of Queen's University, and Wenyu Wang of Indiana University, has provided more evidence to show how the creditor-shareholder conflict affects the resolution of financial distress.

The researchers theorized that reducing shareholder-creditor conflict can ease the path to cost-effective restructuring for distressed firms. They found that when financial institutions hold simultaneous debt-equity holdings in firms, those firms are more likely to restructure out of court than file for bankruptcy. This effect is stronger when loans are over-secured and when bankruptcy costs are higher.

Their paper, entitled "Simultaneous Debt-Equity Holdings and the Resolution of Financial Distress," provides the first empirical evidence of how debt-equity simultaneous holdings affect the resolution of financial distress.

To conduct their research, they assembled one of the largest sets of debt restructurings (494 out-of-court restructurings and 502 bankruptcy filings) between 2000 and 2014. They also assembled a sample of simultaneous holders of loan and equity, and used a labor-intensive matching procedure to identify loan-equity simultaneous holders of distressed firms, and the extent of their holdings. They found that 22 percent of the distressed firms had at least one simultaneous holder.

Analyzing their data, they found that the presence of simultaneous loan-equity holdings is associated with a 35.6 percentage-point increase in the probability of out-of-court restructuring, compared with the 49.6 percent unconditional probability of such restructurings.

"Basically, we find that when a firm has simultaneous debt and equity among its holdings, it's easier for them to resolve their distress by negotiation, instead of filing for bankruptcy," Diep-Nguyen said.

They also found that the presence of simultaneous holdings is associated with 28 to 35 percentage-point higher likelihood of out-of-court restructuring for highly collateralized secured debts compared with those with less collateralized debts. This finding supports their theory that the effect of simultaneous holdings on out-of-court resolutions is stronger when the shareholder-creditor conflict is more severe.

Among their other results, the researchers found evidence that simultaneous holders' incentives to avoid bankruptcy are stronger when the expected bankruptcy costs are high and benefits are low.

"Out-of-court restructurings or workouts are much more cost-effective," Diep-Nguyen said. "Once a firm has to go through bankruptcy, they have to pay a lot of legal fees and they cannot operate normally, so there's a lot of lost revenue."

The researchers also looked at stock performance and found evidence that simultaneous holders can capitalize on cost savings in a successful workout. Firms with simultaneous holdings earn higher stock returns by about 4.5 percentage points each month in the resolution process than firms without.

"If they cannot capitalize on those cost savings, there will no incentive for them to avoid bankruptcy," Diep-Nguyen said. "The stock market is a channel where they can capitalize on those cost savings."

When Real-World Problems Have Too Many Potential Outcomes, Universal Randomized Algorithms May Help Decision Makers

By Melvin Durai

Imagine a scenario where the federal government is preparing a plan to manage Covid-19 infections in Indiana and Illinois, and classifies each state as either "improving" or "worsening." For these two states, the number of possible combinations is four. But if Michigan is added to the plan, the number doubles to eight. And if all 50 states are considered, the number of combinations is astronomical.

"Each time you add a component of interest to the model (like another U.S. state), the complexity of the problem doubles," said William B. Haskell, assistant professor in Krannert School of Management. "In this way, we see exponential growth in the complexity of the problem."

When the state space (or set of all possible configurations) of a problem is finite, a wide variety of approximate dynamic programming (ADP) and reinforcement learning (RL) algorithms exist for Markov decision processes (MDP), which provide a framework for sequential decision making under uncertainty.

But when the state space is very large or continuous, as in the Covid-19 example or many other real-world problems, figuring out which ADP and RL algorithms to use is more challenging.

Many ADP methods require problem expertise to use effectively, Haskell said, citing the example of the ADP used to optimize transportation networks of several major U.S. corporations.

"Supply chain managers and other subject matter experts can propose good strategies for this class of problems because they have been working within the system of interest for years," he said.

Haskell and his collaborators, Rahul Jain and Hiteshi Sharma of University of Southern California and Pengqian Yu of IBM Research Singapore, sought to find ADP methods that do not require subject matter experts to get good solutions.

"We want methods that are sufficiently universal that they can produce good strategies using only data readily available about the problem," Haskell said.

In their paper entitled "A Universal Empirical Dynamic Programming Algorithm for Continuous State MDPs," published in January 2020 in IEEE Transactions on Automatic Control, the researchers propose randomized-function fitting ADP algorithms that are universally applicable for continuous state-space MDPs with finite action spaces.

Explaining these empirical value learning (EVL) algorithms, Haskell noted that one way to solve a dynamic programming or "planning" problem is to compute its "value function."

"The value function tells you the value of every possible state that you might enter, so you can use the value function to try to stay in higher-value states," he said.

Referring to the aforementioned Covid-19 planning example, he said the value function should help direct resources in a way that stabilizes the entire country as quickly as possible.

"Our idea, based on 'randomized value functions,' is needed because it is usually hard to compute the value function exactly," Haskell said. "So we randomly generate a batch of candidate value functions, and then construct our estimate of the true value function from this batch. This idea of using randomly generated or simulated data is a ubiquitous tool to get good solutions to problems that are otherwise impossible to solve exactly."

A secondary contribution of their paper pertains to the random operator framework. As Haskell explained, many algorithms in machine learning, optimization, and dynamic programming can be understood as "operators." Operators are mappings from a class of objects to the same class of objects. For example, dynamic programming problems can be solved by using an operator that accepts a guess about what the value function should be and then returns an improved guess.

"Our algorithm is based on random sampling and simulation, so the operator in our case is 'random,'" Haskell said "As part of this paper, in addition to the algorithm itself, we have developed a new way to analyze this class of random operators and to show that they return a 'good' solution to the problem at hand with high probability."

How Age Diversity Can Help Organizations Perform Better

By Melvin Durai

A woman in her 60s works in a busy office, seated near a recent college graduate in his 20s, with people of varying ages occupying desks nearby. That's an increasingly common sight in America, part of a global trend of rising workplace age diversity. The percentage of American workers aged 55 or older rose from 13 percent in 2000 to 24 percent in 2019, according to the U.S. Bureau of Labor Statistics.

Source:
Yixuan Li

Scholars have studied the relationship between age diversity and organizational performance, suggesting that firms gain a competitive edge through a diverse workforce. But exactly how this diversity may lead to potential benefits has not been clear.

Yixuan Li, assistant professor in Krannert School of Management, has collaborated on a study that helps uncover the mechanisms through which age diversity affects organizational performance.

"We know very little about whether age diversity can actually bring value to the workplace and what the value-creating mechanisms would be," Li said. "That's the main goal of the study — to take the first step toward understanding the potential value-creating mechanisms of age diversity and whether organizations can actually manage age diversity to benefit their performance."

Adopting an intellectual capital perspective, Li and her co-authors examined how age diversity impacts an organization's performance by shaping its human and social capital. Human capital comprises the knowledge, skills and abilities of employees, while social capital refers to the knowledge gained though their social connections.

The researchers found that younger and older employees contribute to human and social capital in complementary ways, allowing their organization to benefit from a broader range of knowledge, skills, abilities and social connections.

They also found that organizations with functional diversity and age-inclusive management are better suited to capitalize on age-diverse workforces.

To conduct their study, the researchers analyzed two-wave data from almost 3,900 workplaces, collected through an annual workplace survey by the Alexandria, Va.-based Society for Human Resource Management (SHRM). Their results appear in a paper entitled "Leveraging Age Diversity for Organizational Performance: An Intellectual Capital Perspective," published online last spring in the Journal of Applied Psychology.

"We hope this provides a more rigorous and comprehensive evidence about how age diversity may actually benefit organizations in general across different industries," said Li, who collaborated with Yaping Gong of Hong Kong University of Science and Technology, Anne Burmeister of Erasmus University Rotterdam, Mo Wang of University of Florida, Valeria Alterman of University of Miami, and Alexander Alonso and Samuel Robinson of SHRM.

The researchers described how age diversity enhances human capital, as younger and older employees complement each other with their knowledge, skills and abilities. For example, older employees have valuable job-specific and social knowledge, while younger employees have more up-to-date scientific and technical knowledge from outside sources. Older employees often have better political and social skills, whereas younger employees are more adept at using advanced learning tools to access new information.

Regarding abilities, older employees have higher crystallized intellectual abilities and can also use their experience to solve new problems. Younger employees, on the other hand, have fluid intellectual abilities and more flexible knowledge structures, allowing them to operate effectively in dynamic, ambiguous and complex environments.

Age diversity also enhances social capital, with older employees developing a close and targeted network, while younger employees focus on information-gathering and develop a broad network.

"This also creates the complementarity," Li said. "If we have people of different ages in the organization, then we can have various types of social connections that can help the organization do better."

The impact of age diversity is amplified in workplaces with high functional diversity, the researchers found.

"This is actually an indication of whether the workplace needs very complicated and diverse knowledge to deal with external contingencies," Li said.

Age-inclusive management also amplifies the impact of age diversity by integrating employees of all ages into the fabric of the organization. Quite notably, the researchers found that the effect of age diversity on social capital became negative when age-inclusive management was low.

"If the organization fails to use age-inclusive management in the workplace, it's going to actually increase age stereotype and discriminatory climate in the organization," Li said. "So in this case, age diversity may actually hurt social capital and people might not be willing to share information obtained from external network connections with the organization. That indicates that age-inclusive management is really important. If organizations fail to use age-inclusive practices, perhaps those older workers and younger workers are not willing to share knowledge between them and they're not willing to share network-based external information with the organization."

The research shows the importance of recruiting and retaining an age-diverse workforce, implementing age-inclusive management, and treating employees of all ages in a fair, nondiscriminatory way. "It's important to develop this inclusive climate, so that people treat each other with respect and people's perspectives are valued," Li said.

Doctors Competing Against Each Other: What's Good for Patients May Not Be Good to On-Demand Healthcare Platforms

By Melvin Durai

Just as you can get a ride through Uber and book a room through Airbnb, you can use an on-demand service platform to make an appointment with a doctor. The doctor will see you through a computer, tablet or smartphone.

Such remote consultations help extend healthcare resources from cities to rural areas and may also keep patients from making unnecessary visits to hospitals. Companies offering telemedicine through online platforms include Amwell and Doctor On Demand in America and ChunYu Doctor in China.

Doctors may provide service through these healthcare platforms in their spare time, while working full-time at clinics or hospitals — or they may even offer their services primarily through the platform. Most are independent contractors on the platforms, competing against other doctors providing the same type of care.

But while competition among doctors may benefit patients, too much competition may hurt the platform.

"If you allow doctors to compete very fiercely with each other, they will hurt the revenue of the platform because doctors are going to compete on price," said Yixuan Liu, assistant professor in Purdue University's Krannert School of Management. "They are going to lower the price and the platform will not earn enough revenue."

Liu and her collaborators, Xiaofang Wang of Renmin University of China, and Stephen Gilbert and Guoming Lai of McCombs School of Business at The University of Texas at Austin, used a mathematical model to study competition and participation in on-demand healthcare platforms, which have been growing in popularity, particularly during the coronavirus pandemic.

"Our study provides some useful insights for the operations of such types of healthcare platforms," Liu said. "It's an emerging business model, but little has been studied about the strategies of the three parties involved in such platforms."

As they describe in a paper entitled "Pricing, Quality and Competition at On-Demand Healthcare Service Platform," the researchers used their model to derive a distributive equilibrium that could be divided into three regions according to the commission rates set by the platform.

In the region with high commission rates, a small number of doctors participate, not enough to serve all the patients. In this region, when commission rates are decreased, doctor participation and market coverage increase, while service quality and price remain the same.

"Very intuitively, if the platform charges a high commission rate, physicians are not that willing to join the platform," Liu said. "So there will be a smaller number of doctors joining the platform. If there are a very small number of doctors on the platform, they won't compete with each other. There is ample demand for them to serve."

In the region with intermediate commission rates, doctor participation is high enough to serve all the patients, and doctors compete on service quality but not price. In this region, a decrease in commission rates increases doctor participation and service quality, as well as price.

The third region has low commission rates and even more doctors participate, competing on both service and price. A further decrease in commission rates increases doctor participation and service quality, but lowers the price.

"It's very critical for the platform to manage competition among doctors," Liu said. "The platform can use various tools, such as controlling the service price or lowering the commission rate to control the competition. That is the main implication of our study."

Using the equilibrium framework, the researchers also investigated the impact of two types of healthcare platforms: one that controls commission rates only and another that controls both commission rates and service prices. They found that when a platform sets commission rates only, it may attempt to set a higher rate to curb competition. But when it controls both commission rates and service prices, it may lower the commission rate while setting a high service price. This increases doctor participation and service quality, benefiting both doctors and the platform without hurting patients.

The researchers also explored what happens when patients have different needs or when the service quality of doctors varies significantly. They found that the platform is most profitable when patient heterogeneity is either small or large – not in between. They also found that the platform generally makes more profit when doctors are more heterogeneous.

Collaborating with a major platform in China, the researchers have also worked on an empirical study to investigate the healthcare platform's design strategy.

Acquisitions or Alliances: Why Companies Persistently Choose One Path over Another

By Melvin *Durai*

Cisco Systems, the American networking company, made its first acquisition in 1993 and has since completed more than 200 acquisitions. While the company has also participated in alliances and joint ventures, it has fueled its growth primarily through acquisitions. The computer firms HP and Dell, on the other hand, have over the years leaned more heavily on alliances and joint ventures for growth.

Companies need to continually reconfigure their resource bases to remain competitive, but how they do this varies significantly. As Cisco, HP and Dell exemplify, some favor a transactional vehicle (acquisitions and divestitures), while others prefer a relational vehicle (alliances and joint ventures).

Luis A. Rios, assistant professor in Krannert School of Management, has collaborated on a longitudinal study that shows firms' persistence in their choice of vehicle, as well as the role of the environment in determining the outcomes of these choices. Firms that persistently choose the relational mode outperform those favoring transactional vehicles during stable economic periods, but the reverse holds true in recessionary periods.

"We were happily surprised to find that the interaction between type of vehicle and external environment was in fact significant," Rios said. "There is almost no prior work doing this, so we did not know what to expect."

Rios and his co-authors, Deepak Jena of Indian School of Business and Patia J. McGrath of University of North Carolina, explored how a firm's resource management history — its accumulation and curation of particular resources — helps determine its future vehicle choices and performance.

In a paper entitled "Persistence in Resource Vehicle Choice and Firm Performance: The Moderating Role of the Environment," the researchers share the results of their cross-industry study of more than 1,300 firms over eight years, providing large-scale evidence of persistence — how firms' choice of vehicle for their resource base reconfiguration endures over a long period.

They argue that while persistence can hinder a firm's ability to adapt, it may also keep the firm from overreacting to market signals, which could potentially hurt more nimble firms.

"There are well-documented 'bandwagon effects' where firms mimic other firms' strategies, especially if they appear successful to the rivals," Rios said. "A typical example is merger waves. So if the market seems to be saying that M&A (mergers and acquisitions) is the best strategy (for example, lots of liquidity or high valuations in the market), a firm that is slower or more cautious in jumping into acquisitions may avoid a winners-curse problem or engaging in unnecessary bidding for targets. Path dependency in alliances accomplishes this. On the flip side, as a hot M&A market cools, if most firms exit quickly, those that remain longer may benefit from bargains."

The researchers suggest that this persistence is entirely rational and can sometimes be beneficial. Switching modes not only means that a firm must invest in different resources and capabilities to support the alternative vehicle, it loses the opportunity to leverage the competitive edge that it has gained through its current resource base.

"This is no different from a car maker that has invested heavily in fossil fuel technology," Rios said. "From an economic perspective, it would take a very large increase in profits to switch to electric. Similarly, a firm that has accumulated resources and capabilities as an acquirer, for example, may be better off staying the course even if alliances were the best option on average."

From a managerial perspective, the research highlights the importance of having a clear understanding of what sets a firm apart from the average, Rios said.

"This sort of self-knowledge is something that is often not done deliberately," he said. "While some firms develop a reputation (and identity) as 'acquirers' (like Cisco), most firms may be hard-pressed to know their own 'type.'"

Statistical Inference for Stochastic Gradient Descent: A Fully Online approach to Construct Confidence Intervals

By Melvin Durai

The stochastic gradient descent (SGD) algorithm is a widely used optimization method for minimizing an objective function to estimate the parameters of a model. Employed in many machine learning tasks, SGD offers a number of advantages over traditional deterministic optimization methods. It's faster to compute, requires less memory and can converge faster for large datasets. It is particularly well-suited to online settings, such as search queries and transactional data, because it needs only one sample at a time to update the parameters.

Source:
Yichen Zhang

While most research related to SGD focuses on the convergence of the objective function or the error of the obtained solution, Yichen Zhang, assistant professor in Krannert School of Management, has collaborated with three other researchers to investigate the problem of statistical inference for SGD and develop a fully online approach to construct asymptotically exact confidence intervals for large-scale data sets.

Statistical inference (constructing confidence intervals) for each coordinate of the minimizer of a convex objective function has been largely unexplored. Confidence levels are used in statistics to quantify the uncertainty in an estimate of model parameters.

Zhang, collaborating with Xi Chen of New York University, Jason D. Lee of University of Southern California, and Xin T. Tong of National University of Singapore, investigated the problem of statistical inference of true model parameters based on SGD when the population loss function is strongly convex and satisfies certain smoothness conditions.

In their paper entitled "Statistical Inference for Model Parameters in Stochastic Gradient Descent," published in February 2020 in The Annals of Statistics, they propose computationally efficient methods to conduct inference for each coordinate of the minimizer, giving an estimate as well as a confidence interval, such as whether there's a 95 percent probability that it will fall within a particular range.

Zhang cites the example of a pharmaceutical company developing a new drug and needing to make a hypothesis test to determine whether the drug is valid to treat a specific illness. A confidence interval allows them to state that the drug will, for example, have a positive effect in 95 percent of the cases.

In the process of constructing an asymptotically valid confidence interval, the researchers needed to construct a covariance estimator that could be used online. The standard estimator required more data storage, which undercut one of the advantages of SGD.

"Traditional methods are usually only able to provide the covariance matrix after we collect all the data, which means we collect the results from a number of patients or experiment targets," Zhang said. "Our method aims to construct such estimators in an online fashion."

The researchers proposed two methods that provide asymptomatically exact confidence intervals: a plug-in estimator and a batch-means estimator. Each offers advantages, depending on tasks and computing resources. The plug-in estimator has a faster convergence rate, but when computing resources are limited, the batch-means estimator may be preferable.

While those two methods consider fixed-dimension setups, the researchers also tried to determine how to construct confidence intervals in high-dimensional setups.

In the drug-testing example, the dimension corresponds to the number of drugs being tested and the sample size corresponds to the number of patients. While classic statistics usually considers dimensions that are smaller than sample sizes, modern statistics considers larger dimensions.

"That means we are trying to make some valid claim when the treatment or index we are interested in is much more than the patients we have," Zhang said.

To address these situations, the researchers used a recently proposed RADAR algorithm, a variant of SGD, to construct a debiased estimator of each regression coefficient that is asymptotically normal.

"The two types of problems, fixed dimension versus high dimension, have significantly different properties," Zhang said. "That's why we deal with them in different sections of our paper."

A Better Way to Measure the Impact of Customer Service Satisfaction on Loyalty

By Melvin Durai

You just got off the phone with a customer service representative at a credit card company. The rep was friendly and helpful, making you feel good about your relationship with the company. A minute later, your phone rings. It's an automated recording asking if you'd like to participate in a survey about your customer service experience.

Such surveys are an increasingly common way for companies to measure customer service satisfaction. They want to find out if your customer service experience will help them retain you as a customer.

"They try to use all kinds of survey methods to measure the quality of service, but there are some inherent problems in measuring the impact of service quality on customer loyalty if you use the survey method," said Dr. Guofang Huang, assistant professor of marketing in Purdue University's Krannert School of Management.

The survey method typically introduces multiple sources of bias that distort the relationship between service satisfaction and loyalty.

Huang and his co-researcher, K. Sudhir of Yale School of Management, have proposed an instrumental-variable (IV) approach that addresses these sources of bias and more accurately estimates the effect of customer service satisfaction on customer loyalty.

The researchers used the IV approach to analyze survey data from a large credit card-issuing company, sharing their results in a paper entitled "The Causal Effect of Service Satisfaction on Customer Loyalty," forthcoming in Management Science. Their findings suggest that many companies may have "grossly underestimated" the return on investment in customer service.

This underestimation can be attributed to the multiple sources of bias that survey methods introduce, including common methods, attenuation and omitted variables.

Common methods bias may arise, for example, when two customers who have received identical service are asked to provide ratings on a scale of 1 to 10 for satisfaction and willingness to recommend (i.e. loyalty intent). One customer may provide more generous ratings than the other, despite receiving the same quality of service, which creates an inflated correlation between service quality and loyalty.

Attenuation bias is introduced because the survey does not provide a perfect measure of a customer's views. If a customer wants to give a rating of "8.5," the only choices are 8 and 9, and the customer is forced to pick a rating that's not perfectly accurate.

Common methods bias tends to cause overestimation of the relationship between service satisfaction and loyalty, while attenuation bias tends to cause underestimation of this relationship.

If only one of these biases exists, the estimate may still be informative, Huang said, but when they're both present, the picture is muddled considerably.

"Once you have these two sources of biases, then companies don't really know where they are," he said.

Another source of bias, omitted variables, occurs because the surveys don't measure every factor that might affect a customer's loyalty. Even if a credit card company's customer service is poor, a customer may remain loyal because, for example, the company offers a very low APR (annual percentage rate).

"If you combine all these sources of bias together, you really need a technique to help you get rid of them," Huang said.

The IV approach does this by taking advantage of randomness – the random way in which service employees are assigned to customers. When a service rep has finished helping one customer, the rep helps the next customer waiting in the queue, and the availability of the reps is independent of the customers in the queue.

Using survey data that identified which customer service representative assisted each customer, the researchers were able to estimate the skill level of each rep. Because the reps were assigned to customers randomly and independent of various factors (including a customer's credit score, credit limit and tenure with the company), the researchers were able to isolate the impact of service quality (and thus satisfaction) on loyalty.

"Whether you were handled by a good or bad service agent has nothing to do with, say, whether you are satisfied with your APR or not," Huang said. "That part of your reporting of satisfaction has nothing to do with other factors."

The researchers measured loyalty in two ways: whether a customer was willing to recommend the company to a friend, as well as whether the customer actually stayed with the company.

They found that not only have companies underestimated the impact of service satisfaction on loyalty, this impact is even greater for more difficult or important calls.

The IV approach can be applied to many other settings, such as banks and hospitals, where service quality is random, Huang said.

"If you walk into urgent care, which doctor serves you really depends on the availability of the doctors," he said. "We're definitely interested in how doctors' service affects our satisfaction with medical services, or our loyalty. If their healthcare service quality is higher, how does that affect our loyalty to a healthcare organization?"

PURDUE UNIVERSITY®

Krannert School of Management

Writers:
Melvin Durai

Faculty:
Diana Choi
Ha Diep-Nguyen
William Haskell
Guofang Huang
Yixuan Li
Yixuan Liu
Luis Rios
Yichen Zhang

Krannert School of Management Department Head Office:

Ananth Iyer, Department Head

Angie Arrigo,
Senior Administrative Assistant
Christa Jackson
Dean's Office Administrative Assistant