MOODY'S ANALYTICS

ANALYSIS

Can Vehicle Residual Forecasts Lift Auto Credit Scores?

Prepared by

Tony Hughes
Tony.Hughes@moodys.com
Managing Director

Contact Us

Email
help@economy.com

U.S./Canada
+1.866.275.3266

EMEA
+44.20.7772.5454 (London)
+420.224.222.929 (Prague)

Asia/Pacific
+852.3551.3077

All Others
+1.610.235.5299

Web
www.economy.com
www.moodysanalytics.com

Introduction

In all collateralized loan forms, the value of the underlying asset is central in the determination of loss given default. When it comes to the probability of default, meanwhile, lending products are more varied in their treatment of collateral value. In mortgages, for example, a loan-level PD model that does not contain the origination loan-to-value ratio or ongoing changes in house prices would not pass any satisfactory validation check. In the world of autos, in contrast, projections of vehicle value traditionally play a much more subdued role. In most stress-testing applications, forecasts of used-car price indexes are typically employed to track changes in collateral value over time. In scoring applications, meanwhile, it is far rarer to see models that take explicit account of expected future vehicle value.

We seek to explore this issue in this paper.

Can Vehicle Residual Forecasts Lift Auto Credit Scores?

BY TONY HUGHES

In all collateralized loan forms, the value of the underlying asset is central in the determination of loss given default. When it comes to the probability of default, meanwhile, lending products are more varied in their treatment of collateral value. In mortgages, for example, a loan-level PD model that does not contain the origination loan-to-value ratio or ongoing changes in house prices would not pass any satisfactory validation check. In the world of autos, in contrast, projections of vehicle value traditionally play a much more subdued role. In most stress-testing applications, forecasts of used-car price indexes are typically employed to track changes in collateral value over time. In scoring applications, meanwhile, it is far rarer to see models that take explicit account of expected future vehicle value.

We seek to explore this issue in this paper.

Suppose we have two potential buyers of new vehicles. Each has the same credit score and both individuals are buying a $20,000 car with a $5,000 down payment. They both earn the same amount. The key difference is that one is intending to buy a VW Golf, while the other has her sights set on a Jeep Renegade. Can we utilize this additional piece of information to better measure the riskiness of each borrower? Will knowledge of vehicle characteristics help us to better rank the individuals in the data from most to least likely to default?

It is reasonable to believe that such information will provide lift. It is well known, for example, that certain vehicles consistently retain their value better than their rivals. A smart shopper, realizing that the total cost of ownership includes depreciation, will tend to seek out vehicles that are known to retain their value better. Although cars have generally improved in build quality and expected life in recent years, some brands enjoy a reputation for providing many years of reliable driving and others are derided for their unreliability. All borrowers sustain negative financial shocks at various times. If they do, it makes sense that they would fight harder to defend vehicles with better potential long-term motoring prospects from the threat of repossession.

On a more humanistic level, it is often said that the car we own defines our personality. The central premise of this paper is that vehicle choice could provide an additional indication of creditworthiness even after controlling for the financial terms of the deal and the credit score of the borrower. Could it be, for example, that sensible, safety-conscious individuals buy Toyota Camrys and pay their bills on time, while Chevy Malibu owners are more fun-loving and carefree and thus more likely to default?

To address questions such as these, we source data on credit performance from EDGAR Online covering all U.S. auto loans that have been packaged and sold as asset-backed securities. Under the Securities and Exchange Commission Regulation AB II, bond issuers are required to publish detailed loan-level data on the performance of the underlying assets, including information on delinquencies, defaults, and credit losses that stem from the loans. The database also contains the FICO score at origination, interest rates and loan terms, and information concerning the vehicles serving as collateral. The data have been available in this form only since January 2017, though this still gives us 15.7 million consistent observations through November 2017, the snapshot used for the analysis conducted here.

We combine data on credit performance with vehicle information drawn from Auction.net and related sources. We have used this database previously to produce forecasts of vehicle value, including price-to-manufacturer's suggested retail price ratios under a range of economic scenarios, including a baseline forecast of residual value at all future time points.

Our central question concerns whether these forecasts are useful for assessing the creditworthiness of borrowers. To this end, we build credit scores that utilize residual price forecasts and static vehicle information and compare the performance of these models with traditional scoring techniques that are blind to the nature of the vehicle backing the loan.

In the next section we will describe the methodology we employed to do this before diving into the empirical results of our exploration.

Methodology

Vehicle information in the ABS data is somewhat limited in the sense that only make, model and model year are recorded. We also learn whether the vehicle was most recently purchased in a new or used condition. We do not gain any information regarding the 10-digit vehicle identification number, the trim of the vehicle, or its initial mileage if purchased used. We join the two databases using text strings available both in the ABS data and in Auction.net. We throw out any observations that cannot be matched either because of typographical errors in the description of the vehicle (though we cover easily addressed variations in the matching process) or because of a lack of vehicle information in general. Data from Auction.net are, for all intents and purposes, universal so any that are missing are due to problems with the vehicle descriptions provided by the ABS data. Our hit rate was well above 90%. We suspect that the prevalence of typos is uncorrelated with any results presented here.

To the best of our ability, we score the loans in a completely out-of-time, out-of-sample manner. Our aim is not to build origination credit scores, per se, but to predict monthly probability of default at all points in the life cycle. We compute our scoring regressions using data up until August 2017 and develop forecasts of vehicle value using models estimated in the same month. We then assess the performance of the scoring models during the final three months of the observation window, September through November 2017. We have about 15.7 million individual observations of loan performance in total, of which 7.6 million are populated with vehicle information and fall in the estimation window.

Our baseline model is designed to measure the effectiveness of the FICO credit score in predicting the probability of severe delinquency. We define this event as being 60 or more days past due during the period of interest. Our baseline model also includes a life-cycle function developed using a liberally specified piecewise linear spline function. We model interactions between the life cycle curve and the credit score so that accounts of different quality are allowed to season at different rates. We estimate all models using logistic regression.

The ABS data are made up of deals originated by a number of organizations that are active in the ABS market. This list of organizations is not completely representative—it is dominated by large subprime lenders and most of the major captive finance companies. Traditional banks, which tend to keep auto loans on their balance sheets, are underrepresented in the data. We control for differing standards among the issuers by including a full set of relevant dummy variables. In reality, differences in the supply of credit available to various vehicle makes may be a determining feature of observed credit performance. Here we estimate the models by reduced form and attempt to control for variations in supply propensity by using issuer dummies.

The first generalization we consider takes the baseline model and appends a variable that represents a baseline forecast of the vehicle's price-to-MSRP ratio drawn from Moody's AutoCycle. As mentioned previously, this variable is an in-sample, model-based prediction of the vehicle's fundamental value during the estimation period and a true out-of-time, ex ante prediction of the vehicle's value during the holdout sample. In this model, we also include an interaction between the forecast price-to-MSRP and the observed origination FICO score.

We then introduce additional static vehicle information into the baseline model. In this iteration we include variables that directly define the vehicle—a dummy for those sold in a used condition, the vehicle's age at the time of origination, dummies for different manufacturers, and dummies for different vehicle segments using National Automobile Dealers Association's closely followed classification system.

The final iteration involves combining all the elements described in previous paragraphs: the FICO score, life cycle, price-to-MSRP forecast, and static vehicle-specific information. All interaction terms, as defined above, are also included in this specification.

To assess the performance of the model in the holdout sample, our primary focus is on the ability of the models to rank-order individual accounts from most to least likely to default. To this end, we use the estimated models to predict default probabilities and then calculate an associated Kolmogorov-Smirnov statistic for each score. This statistic measures the maximum difference between the cumulative distribution of scores for defaulted accounts and the cumulative distribution of scores for clean accounts. A high KS statistic is the ultimate aim. This statistic is used widely across the industry in the assessment of "lift" provided by various credit scores.

Results

The baseline model, based only on knowledge of the issuer and of the FICO score, achieves a 59.61% KS statistic (see Table 1). By comparison, a model that considers vehicle information only in isolation (while controlling for issuer variation) scores below 57%. The underlying model achieves a pseudo-R2 of 15.9%, all variables are of the anticipated sign, and issuer dummies are jointly highly significant. The FICO score, in isolation, provides a huge amount of information about likely creditworthiness. This result is hardly surprising.

Our primary goal in this paper is to provide evidence that forecasts of future vehicle residual value can be used to better assess probability of default. When we consider the model that appends credit scores with out-of-time forecasts of the underlying vehicle's value, we find that the KS statistic ticks slightly higher, to 59.79%, suggesting that combining standard credit scores with AutoCycle forecasts provides slightly better credit separation than using the raw score on its own.

Consider this result carefully. We are not saying that knowing the future value of the car would enable better credit scoring, we are instead suggesting that accessing a carefully built out-of-time forecast of the vehicle's future value allows creditworthiness to

Table 1: Kolmogorov-Smirnov Statistics Calculated for PD Scores of Varying Complexity

Days past due	FICO score + life cycle + issuer dummies	FICO score + life cycle + issuer dummies + residual forecasts	FICO score + life cycle + issuer dummies + static vehicle characteristics	FICO score + life cycle + issuer dummies + residual forecasts + static characteristics
60+	59.61	59.79	60	60.07
0+	44.95	45.08	44.99	45.2

Sources: NADA, EDGAR Online, Moody's Analytics

be better assessed. The lift is only slight, but it remains remarkable that a model-based prediction, itself prone to error, can be used to better rank the future creditworthiness of vehicle owners.

When we turn to the third formulation that appends the baseline model with known static vehicle characteristics, we find that the KS statistic rises more significantly, to 60%. The model in use here is much more complex than the baseline specification. Bear in mind, though, that added complexity normally detracts from the ability of any model to successfully make ex ante predictions. The fact that we were able to derive a greater amount of separation in such a setting implies that vehicle information really does promote more effective credit scoring than using an off-the-shelf commercial score in isolation. In this case, the vehicle attributes are known with certainty, whereas model-based residual forecasts are merely statistical estimates. One suspects that if future resale value were known at the time of scoring, it would be highly indicative of future credit performance.

We get the best separation, though, when we combine known vehicle information with the AutoCycle forecasts of future price. In this case, the KS statistic rises to 60.07%. This means that, relative to the baseline model, we have gained 0.46 point in terms of KS separation. The residual price forecast clearly provides an additional signal that is not captured by static vehicle characteristics or by the raw credit score.

We tried a number of additions to these models in a bid to boost the KS of the preferred, most general specification described here. Our primary concern was that we were not controlling for repayment burden, nor for the size of initial down payments offered by borrowers in securing their loans. Interestingly, neither of these additional factors lifted the measured KS of the resultant scores. Adding vehicle LTV yielded a KS of 59.7%—a slight rise relative to the baseline model but well below the number derived from the preferred model. Adding the payment-to-income ratio causes the KS to backslide a little more, to 59.67%. These results underscore the notion that gains observed in our preferred model are not simply the result of data mining.

When we instead consider the prevalence of shorter-term delinquency—which is often driven by tardy payments and is far noisier and less consequential than long-term failure to pay—we find that price forecasts provide relatively more lift than static vehicle characteristics. The margins here are thinner, though it is still instructive that vehicle information lifts the ex ante prediction accuracy of standard credit scores regardless of the seriousness of the delinquency considered. When both sources of collateral-related information are applied to these shorter-term delinquency rates, we find a rise in KS of 0.25 percentage point compared with the baseline model.

How economically significant is vehicle information?

In the previous section, we identified small but statistically significant gains in rank-ordering effectiveness for credit scores that incorporate vehicle information. An interesting question to now ask is how the differences between vehicles identified in the models translate into scores and associated interest rates.

Using the baseline model, we can construct curves that represent the relationship between raw FICO score and probability of default. These curves will vary by issuer and by the age of the loans considered, so in this analysis we will isolate loans originated by a specific subprime lender that were 6 months old in November 2017. An individual with a 600 FICO score has about a 2% probability of being 60+ days delinquent given these conditions. The mapping curve is included in the chart.

We then calculate PDs for specific vehicles using our preferred model that incorporates both price forecasts and static vehicle information. The chart maps estimated PDs from this model into an approximate impact measured in FICO score points.

The most popular vehicle in the U.S.—the Ford F-150 pickup truck—also happens to be one of the best-performing vehicles in terms of observed underlying default rates, according to our model. To put this into context,

Credit Score to PD Mapping Curve

FICO (x-axis), 60+ days past due estimated delinq. rate (y-axis)

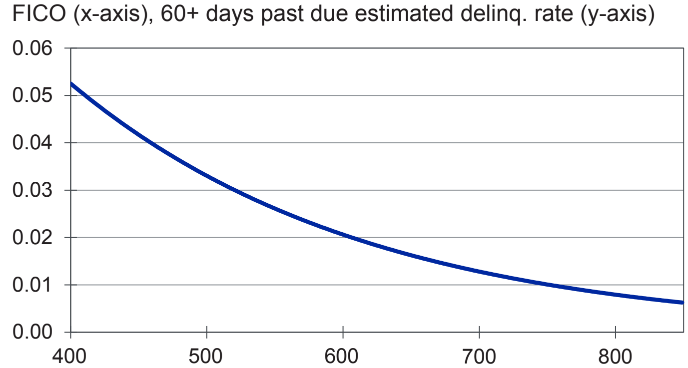

Sources: NADA, EDGAR Online, Moody's Analytics

a driver of a 2014 F-150 with a 600 FICO score was predicted to have a November default probability of 1.3%, well below the 2% prediction for the driver of a generic vehicle. Following our mapping process, this PD translates into a FICO score of 696.

On the flip side, the 2012 Nissan Leaf is perhaps the worst-performing vehicle of late in terms of residual value. A 600-FICO individual driving a Leaf is predicted to fall into the 60-day delinquency bucket 3.6% of the time, which is akin to the behavior of a 490-FICO person driving a generic vehicle. Bear in mind that these are one-month default probabilities conditional on the account surviving to the start of the period in question.

We observe two individuals with precisely the same FICO score, both originated by the same organization in the same time period. According to our vehicle-based score, the two should instead be separated by 206 FICO points! This is a stark finding but not a particularly surprising one. The problems with the resale value of the Leaf have been well documented. The car had a meager 50-mile range when new, and electric cars have been hard-hit on resale value during a period of relatively low oil prices. Americans have been trending toward purchasing large SUVs and pickup trucks, and the Ford F-150 has been a particularly strong seller of late in this category.

In terms of economic significance, we need to consider the terms of loans that will typically be offered to individuals with 490 FICO scores relative to those boasting a 696. According to one source we were able to find, those with a 500-589 FICO will currently expect to find a 15.24% annual percentage rate on a 60-month loan of $20,000. Our score is slightly below this range, which would imply that the actual rate would be slightly north of this figure, perhaps 17% or 18%. A 696, under the same terms, would expect to pay a 4.95% APR. By contrast, a 600 FICO would normally attract a 14% APR, according to the same source.

In other words, choosing the right vehicle translates into a greater than 10-percentage point gap when performance differences are translated into actual borrowing costs.

Vehicle information clearly provides economically significant information that could be used by lenders in the context of credit underwriting.

What vehicle types are associated with better credit performance?

Our preferred model considers residual forecasts as well as static vehicle information. There is also some interest in considering the marginal effects of static vehicle information on subsequent probability of default. In this analysis, for ease of interpretation, we revert to the model that excludes the residual price forecast. The question we ask here is whether there are certain vehicle characteristics that are more indicative of default after controlling for the borrower's standard credit score.

We begin with a consideration of vehicles based on market segment. Table 2 ranks the defined segments in terms of their estimated marginal effect on 60+ day delinquency rates. Those at the top of the list are the least likely to witness a delinquent account after controlling for all the other elements included in the model. The coefficients themselves are somewhat difficult to interpret—the underlying models are, after all, nonlinear logistic regressions. Nevertheless, we include the estimates for reference by tech-savvy readers, along with t-statistics that test whether behavior of the segment in question is significantly distinct from the reference category (in this case, compact utility vehicles).

As mentioned earlier, we have seen some very sharp deviations in residual prices between segments over the past few years as world oil prices have fallen and the American penchant for large trucks and SUVs has come to the forefront. These vehicles, together with large and midsize vans (which are primarily for commercial use), rise to the top in terms of the paucity of associated credit problems. The most surprising aspect of this trend is that midsize and compact SUVs also perform very strongly despite being rather fuel-efficient. It seems that consumers favor the SUV body shape to the point that they are more willing to defend these vehicles against the prospect of repossession.

At the other end of the spectrum we have compact and midsize cars. These vehicles have been hammered on residual value in recent years—this table underscores the fact that future resale value is a critical driving force in the determination of default probability. Large cars and sports cars are the other categories that are found to perform relatively poorly.

Table 2: Marginal Effects on Default Probability by Segment

NADA segment	Coefficient	t-stat
Large van	-0.22	-1.11
Luxury compact utility	-0.18	-1.02
Midsize van	-0.06	-0.92
Luxury large truck	-0.01	-0.09
Compact utility*	0	
Large SUV	0.02	0.25
Midsize utility	0.07	2.12
Luxury sport	0.08	0.47
Luxury midsize utility	0.09	0.65
Large pickup	0.09	2.11
Premium luxury large	0.11	0.5
Midsize pickup	0.14	2.14
Luxury large	0.15	0.84
Upper sport	0.22	2.86
Upper midsize	0.34	5.55
Near luxury	0.35	2.88
Ultra luxury/exotic	0.37	0.94
Sport	0.38	6.63
Large car	0.4	8.34
Intermediate subcompact	0.4	10.33
Intermediate midsize	0.42	13.18
Upper compact	0.44	4.64
Intermediate compact	0.46	14.42
Luxury midsize	0.5	3.84
Entry subcompact	0.66	14.12

*Reference category

Sources: NADA, EDGAR Online, Moody's Analytics

It is important to realize that the shape of this table is very much a function of the times. Were we able to produce the table for a period of high oil prices, it is quite likely that compact cars would be far better represented at the top of this particular league table.

We now turn to vehicle make. Suppose we have two individuals, both with a 750 FICO who both choose to buy a car in the luxury midsize utility segment at the same time. The question here is whether someone who chooses to buy a Lexus is more or less likely to default than someone who chooses to buy a BMW. The ranking of vehicle brands is included in Table 3. The format and interpretation of the numbers are precisely the same as the earlier example.

The ranking of vehicle brands is interesting in a number of ways. The brands that are associated with the lowest default rates include some that are well known for excellent resale properties. Subaru, Toyota/Lexus and Honda/Acura are all historically strong performers in terms of reliability and resale price retention. Ford trucks, as mentioned earlier, include a number of very popular vehicles that are in very high demand from consumers in recent years.

One brand near the top of the list that may be surprising is Smart. These are microcars with very few competitors in the U.S. that are quite practical vehicles for urban drivers. They share some characteristics with the aforementioned Nissan Leaf in that they are among the most fuel-efficient cars on the road. Indeed, Smart cars have been hit hard on residual lately, just like the Leaf. We are controlling for market segment here, so this result simply suggests that among the group of sub-compact vehicles, all of which have performed poorly of late, Smart cars hold up far better than most.

At the other end of the scale, we have a large number of luxury brands. BMW, Porsche, Jaguar and Mercedes-Benz are all located near the bottom of the table. These cars are found to perform relatively poorly in terms of default behavior among the segments in which they are active. Luxury marques are not condemned to lie at the bottom of this table, as evidenced by the fact that Lexus and Acura enjoy low conditional default probabilities. We did, though, identify a stark difference between the Japanese brands and their European rivals in our findings. The other brands in the lower reaches of the table include Mitsubishi, Dodge (trucks and cars alike), Suzuki and Land Rover. A number of these brands have had significant reliability issues.

When additional variables, notably the LTV and payment-to-income ratio, are added to the model, the qualitative nature of the results presented in these tables does not change. As noted earlier, these variables did not lift the rank-ordering ability of the derived credit scores.

To assess the consistency of these rankings over time, we re-estimated our model using the data held over from the final three months of the observation window. The new rankings are included in the right-hand column of Table 3. Some brands, most notably Porsche and Volvo, improved their position, while a handful of others, especially Jeep and Infiniti, fell down the list. Overall, though, the rankings were notably consistent across the two time periods. The Spearman rank correlation between the two series was 0.76;

Table 3: Marginal Effects on Default Probabilities by Make

Make	Coefficient	t-stat	Rank: Jan-Aug	Rank: Sep-Nov
Subaru	-0.3	-1.56	1	3
Lexus	-0.08	-0.59	2	9
Honda	-0.04	-0.31	3	1
Acura*	0		4	6
Smart	0.03	0.1	5	5
Mini	0.03	0.16	6	4
Toyota	0.08	0.59	7	14
Ford truck	0.11	0.75	8	19
Chevrolet	0.14	0.96	9	8
Hyundai	0.15	1.04	10	17
Volvo	0.16	0.78	11	2
Buick	0.16	1.14	12	7
Infiniti	0.17	0.99	13	33
Mazda	0.18	1.15	14	11
Jeep	0.22	1.48	15	28
Cadillac	0.22	1.66	16	12
Volkswagen	0.23	1.52	17	13
Chevrolet truck	0.24	1.66	18	20
Fiat	0.24	1.46	19	16
Kia	0.25	1.75	20	21
GMC light duty	0.28	1.85	21	10
Lincoln	0.29	1.88	22	22
Ford	0.29	2.03	23	18
Chrysler	0.38	2.59	24	25
Nissan	0.4	2.83	25	27
Audi	0.41	2.57	26	30
Dodge	0.41	2.89	27	32
Dodge truck	0.44	3.02	28	35
Mercedes-Benz	0.45	3.47	29	24
BMW	0.51	4.02	30	26
Suzuki	0.54	2.54	31	29
Jaguar	0.58	2.42	32	23
Porsche	0.62	2.36	33	15
Mitsubishi	0.63	4.19	34	31
Land Rover	0.95	4.52	35	34

*Reference category

Sources: NADA, EDGAR Online, Moody's Analytics

the hypothesis of statistical independence between them was rejected categorically. When we next update the results of this research, it will be interesting to see whether the same patterns are observed across different manufacturers.

Vehicles are ranked by many organizations in many different ways—reliability, resale value, consumer satisfaction, safety and so forth. The difference here is that we are trying to distill the underlying propensity of owners of different vehicles to default after a range of lender and borrower characteristics are carefully controlled. We suspect that this tendency to default measure may be an excellent new way to accurately assess customer satisfaction across vehicle brands.

Default, after all, carries severe consequences for car owners. If stressed borrowers are prepared to fight to keep their cars in the wake of financial difficulties, this is perhaps the strongest possible signal of esteem for a particular brand or for a particular vehicle segment.

Conclusion

The results presented here are both curious and tantalizing. The underlying data used for the analysis, though voluminous, have only a short history. In this paper we were able to estimate a model using only eight months of data and then validate the derived scoring methodologies over the subsequent three-month period. Ideally, of course, we would like to evaluate the performance of our technique at a number of points in the business cycle and under a range of fuel price trends. The results are highly promising, but we need more data to be sure.

Despite this, the results strongly suggest that vehicle information, including residual price forecasts, lift the ability of traditional scores to classify borrowers from most- to least-likely to default. These results were achieved using a true out-of-time forecast of residual value, together with a set of easily observed static vehicle characteristics. It is true to say that residual value is a critical determinant of loss given default; the results presented here suggest that it is also a critical determinant of raw default probability.

About the Author

Tony Hughes is a managing director of research at Moody's Analytics. He serves as head of a small group of high-caliber modelers, charged with identifying new business opportunities for the company. Prior to this appointment, he led the Consumer Credit Analytics team for eight years from its inception in 2007. His first role after joining the company in 2003 was as lead economist and head of the Sydney office of the company Moody's Economy.com.

Dr. Hughes helped develop a number of Moody's Analytics products. He proposed the methodology behind CreditCycle and CreditForecast 4.0, developed the pilot version of the Stressed EDF module for CreditEdge, and initiated the construction of the Portfolio Analyzer (ABS) product that provides forecasts and stress scenarios of collateral performance for structured securities worldwide. More recently, he championed and oversaw the development of AutoCycle, a tool that provides forecasts and stress scenarios for used-car prices at the make/model/year level. He has a current development project related to quantifying counterparty network risks that can be applied to the assessment of systemic risk in the financial system.

In the credit field, Dr. Hughes' research has covered all forms of retail lending, large corporate loans, commercial real estate, peer-to-peer, structured finance and the full range of pre-provision net revenue elements. He has conducted innovative research in deposit modeling and in the construction of macroeconomic scenarios for use in stress-testing.

Dr. Hughes has managed a wide variety of large projects for major banks and other lending institutions. In addition, he has published widely, in industry publications such as American Banker, Nikkei, GARP, and the Journal of Structured Finance as well as several papers in peer reviewed academic journals. He obtained his PhD in econometrics from Monash University in Australia in 1997.

About Moody's Analytics

Moody's Analytics provides financial intelligence and analytical tools supporting our clients' growth, efficiency and risk management objectives. The combination of our unparalleled expertise in risk, expansive information resources, and innovative application of technology helps today's business leaders confidently navigate an evolving marketplace. We are recognized for our industry-leading solutions, comprising research, data, software and professional services, assembled to deliver a seamless customer experience. Thousands of organizations worldwide have made us their trusted partner because of our uncompromising commitment to quality, client service, and integrity.

Concise and timely economic research by Moody's Analytics supports firms and policymakers in strategic planning, product and sales forecasting, credit risk and sensitivity management, and investment research. Our economic research publications provide in-depth analysis of the global economy, including the U.S. and all of its state and metropolitan areas, all European countries and their subnational areas, Asia, and the Americas. We track and forecast economic growth and cover specialized topics such as labor markets, housing, consumer spending and credit, output and income, mortgage activity, demographics, central bank behavior, and prices. We also provide real-time monitoring of macroeconomic indicators and analysis on timely topics such as monetary policy and sovereign risk. Our clients include multinational corporations, governments at all levels, central banks, financial regulators, retailers, mutual funds, financial institutions, utilities, residential and commercial real estate firms, insurance companies, and professional investors.

Moody's Analytics added the economic forecasting firm Economy.com to its portfolio in 2005. This unit is based in West Chester PA, a suburb of Philadelphia, with offices in London, Prague and Sydney. More information is available at www.economy.com.

Moody's Analytics is a subsidiary of Moody's Corporation (NYSE: MCO). Further information is available at www.moodysanalytics.com.

About Moody's Corporation

Moody's Analytics is a subsidiary of Moody's Corporation (NYSE: MCO). MCO reported revenue of $4.2 billion in 2017, employs approximately 11,900 people worldwide and maintains a presence in 41 countries. Further information about Moody's Analytics is available at www.moodysanalytics.com.