

METHODOLOGY

Prepared by

Tony Hughes
Tony.Hughes@moodys.com
Managing Director

Samuel W. Malone
Samuel.Malone@moodys.com
Director

Michael Brisson
Michael.Brisson@moodys.com
Economist

Michael Vogan
Michael.Vogan@moodys.com
Economist

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.moodyanalytics.com

Introducing AutoCycle™: Residual Risk Management and Lease Pricing at the VIN Level

Abstract

This white paper lays out AutoCycle™, a model capable of forecasting car prices at the 11-digit vehicle identification number level conditional on a wide variety of macroeconomic scenarios. We demonstrate the core capabilities of our model to capture aging and usage effects and illustrate the material implications for car valuation of different macroeconomic scenarios such as recessions and oil price spikes. Forecasts can be generated for cars of any quality percentile, given car features, as well as for cars of future model years.

In this paper we provide a case study that calculates break-even initial deposit amounts and monthly lease payments for two used cars taken from a recent article on used-car leasing in the Wall Street Journal. We validate our model against a meticulously constructed Challenger model and find that AutoCycle achieves a superior out-of-sample mean error of -0.6% and R-squared of 0.89 on a set of 6.4 million cars whose transactions were recorded during 2015.

The AutoCycle solution is applicable to forecasting the value of lease portfolios, managing residual risk, and pricing individual lease contracts for new or used cars. Our purely quantitative approach allows users to fully validate the model and conduct detailed and transparent sensitivity analyses.

Introducing AutoCycle: Residual Risk Management and Lease Pricing at the VIN Level

BY TONY HUGHES, SAMUEL W. MALONE, MICHAEL BRISSON, AND MICHAEL VOGAN

This white paper lays out AutoCycle™, a model capable of forecasting car prices at the 11-digit vehicle identification number level conditional on a wide variety of macroeconomic scenarios. We demonstrate the core capabilities of our model to capture aging and usage effects and illustrate the material implications for car valuation of different macroeconomic scenarios such as recessions and oil price spikes. Forecasts can be generated for cars of any quality percentile, given car features, as well as for cars of future model years. In addition, we provide a case study that calculates break-even initial deposit amounts and monthly lease payments for two used cars taken from a recent article on used-car leasing in the Wall Street Journal. We validate our model against a meticulously constructed Challenger model and find that AutoCycle achieves a superior out-of-sample mean error of -0.6% and R-squared of 0.89 on a set of 6.4 million cars whose transactions were recorded during 2015. The AutoCycle solution is applicable to forecasting the value of lease portfolios, managing residual risk, and pricing individual lease contracts for new or used cars. Our purely quantitative approach allows users to fully validate the model and conduct detailed and transparent sensitivity analyses.

1. Introduction

Accurate auto residual price forecasts are more important than ever. The market for new cars has, at last, fully recovered the ground lost during the Great Recession, and the industry now looks to settle in for steady growth in line with the outlook for the broader economy. The key risks associated with financing vehicle purchases—be they lease or loan, prime or subprime, fleet or individual, new or used—are invariably realized when cars are sold into the secondary auto market. Anyone with a pecuniary interest in the value of large numbers of vehicles should be keen to sharpen his or her quantitative awareness of the dynamics of such markets.

In this paper, we introduce a new tool for analyzing these dynamics. Our models output forecasts and stressed macroeconomic scenario projections for wholesale used-car prices at an 11-digit VIN level. The models capture differential effects of supply- and demand-side macroeconomic drivers and fuel prices on observed vehicle values. We can accurately differentiate, for example, between two cars of the same type that differ only in their observed mileage. In addition, our model is regional, so heterogeneous macro conditions in different parts of the country can be fully considered. The model projects the likely performance of new vintages of vehicles based strictly on

the past evolution of observed prices within the brand. Further, we are able to differentiate between vehicles that are presented for sale at a variety of quality or condition levels, after controlling for VIN-level vehicle characteristics, the region in which the transaction occurs, as well as the vehicle's observed mileage.

For example, we can project the likely wholesale price of a black 2012 Audi Q5 that has a tan interior, has 42,000 miles on the odometer, is expected to be driven 12,000 miles per year, presents in a condition that is better than 80% of similar vehicles, and will eventually be sold in California. Not only can we provide a baseline forecast for this ve-

hicle, we can also tell you the possible effect on the car's residual value of a dire recession, a surge in oil prices, a tightening of financial conditions, or of all of these calamities occurring simultaneously.

The auto recovery has been driven by two key financing trends—a rapid rise in leasing generally and the development of new, somewhat more creative forms of funding aimed at the lower end of the credit-quality spectrum. Novel forms of finance have taken two distinct paths. On the one, loan terms have been extended to ensure that income-strapped motorists can afford the repayment burden presented by their new or updated vehicles. The paying down of principal occurs at a lower rate than is typical for traditional shorter-term commitments. With borrower equity remaining lower for longer, lenders must remain more vigilant in tracking and predicting the exact financial position of their borrowers, especially those with poorer credit. If defaults occur in longer-term loans, there is a significant risk that the recovered value of the vehicle will fall short of outstanding principal.

The second path is the rapid diffusion of used-vehicle leases. With cars now retaining their utility for longer because of improved manufacturing standards, it makes sense for financiers to offer leases on vehicles that have a few years of history. Used-vehicle leases, however, require a new form of residual calculation to be effective and profitable for the lessor. For brand-new vehicles, yet to be sold into the secondary market, it may be possible or desirable to base future residual value on a subjective assessment of the vehicle. For used cars, however, signals regarding the likely trajectory of a car's value have already been accumulated. For these used-vehicle lease transactions, we can analyze and project existing data and not rely on

out-of-date subjective assessments to forecast a vehicle's market value.

Lessors of new vehicles could benefit greatly from purely quantitative assessments of car values. Most vehicles that are subject to lease agreements are well-established brands with reams of data on likely resale prices. Many auto industry insiders focus heavily on the effect of brand makeovers and redesigns, positing that these make inference from one model year to the next difficult. Sometimes these redesigns can have a profound impact on residual value, but such situations are quite rare. In cases where the redesign is truly revolutionary—the first hybrid or the first car with airbags—it will be difficult to objectively assess the impact of the innovation on future residuals. More commonly, redesigns represent manufacturers' "swimming to keep up with the prevailing current" lest their vehicles lose market share to their competitors.

Some redesigns are also unforced errors. Shoppers may react badly to unnecessary design changes to beloved models. Occasionally manufacturers will accidentally or intentionally switch a model to an adjacent sector and thus lose market traction. In part, these changes in design will be reflected in higher (or lower) sticker prices without really affecting the rate at which depreciation occurs. Ultimately, this is an empirical question, and forecasts from our model can be compared for specific vehicles to show how their prices evolve after future redesigns.

For lessors of new vehicles, who currently rely on partly subjective residual calculations, the pure objectivity of a strictly data-based approach should still bring significant utility to users of the forecasts. Those—such as large banks—managing their businesses under a regulatory spotlight need analytical solutions in which any managerial overlay can

be isolated and separately reported. Models used in the risk-assessment process must be amenable to validation and backtesting, and sensitivity to changes in input variables must be assessable. This requirement holds as much for Challenger models, which are used by banks to keep in-house champion models finely honed and fit for use, as it does to the champion models themselves.

Other, unregulated users will also benefit from these features. We recognize candidly that industry insiders often track brand-specific trends and production data far more closely than we ever would be able to. Nevertheless, in a competitive marketplace such as the auto industry, market forces and other macroeconomic supply- and demand-side drivers will still play a dominant role in determining the behavior of prices for individual vehicles or vehicle segments. We envisage that users will employ our forecasts and stress scenarios as the starting point for appropriate, partially subjective forecasts produced by their own in-house teams.

The rest of this paper proceeds as follows. Section 2 lays out the AutoCycle model methodology as well as the methodology of a Challenger model that we use to validate AutoCycle's performance. Section 3 describes the data. Section 4 lays out the results, with an illustration of aging and usage effects contained in Section 4.1; Section 4.2 is a discussion of macroeconomic scenario-driven price forecasts; Section 4.3 demonstrates our algorithm for generating forecasts as a function of vehicle quality; Section 4.4 presents our solution for forecasting the value of vehicles of future model years; Section 4.5 contains the used-car leasing case study; and the model validation can be found in Section 4.6. Section 5 concludes.

2. Methodology

In this section, we first describe the AutoCycle model, followed by the Challenger model we match up against it in the validation exercise whose results we report in Section 4.6.

AutoCycle

AutoCycle uses a linear model to capture the relationship of residual vehicle values to VIN-level car features and the macroeconomic environment. The dependent variable

for the model is a logit transformation of the vehicle's sale price as a fraction of its manufacturer suggested retail price, or MSRP. A logit transformation restricts the price-to-MSRP forecasts to the interval (0,1). The

independent, or right-hand-side, variables of the model include time-invariant vehicle features, time varying mileage per year and vehicle age variables, month-of-year dummies to capture seasonal effects, and macroeconomic variables.

We use the miles-per-year determinant of car value, rather than a mileage variable, for two reasons: Mileage has a time trend that renders the variable nonstationary in levels, and measuring mileage in this way would conflate overall miles driven with the car's age. Employing mileage per year as the car-usage variable solves both of these problems. Since both age and MPY are included in the model, we can easily reconstruct projections for vehicles with any mileage at any time.

Regarding macro drivers of car value, we include drivers that allow us to capture the divergent behaviors often displayed by different cars during periods of economic stress. Some of these economic variables, such as the unemployment rate and the year-on-year growth rate of disposable income, are meant to capture demand, and others, such as the growth of new-vehicle registrations, are meant to gauge the supply of vehicles in the market during the sale period. A representation of the model, which is a reduced form capturing underlying supply and demand conditions, can be seen in Equation (1) below.

AutoCycle Model Equation

$$(1) \ln(y_{it}/(1 - y_{it})) = \alpha + \beta_1' \overrightarrow{Feat}_{it} + \beta_2' \overrightarrow{TimeFeat}_{it} + \beta_3' \overrightarrow{Macro}_{it} + \beta_4' \overrightarrow{Seg}_{it} \overrightarrow{FeatMacro}_{it} + \beta_5' \Delta U E_{it} \overrightarrow{SubFeat}_{1it} + \beta_6' \Delta Gas_{it} \overrightarrow{SubFeat}_{2it} + \beta_7' \overrightarrow{Debt}_{it} \overrightarrow{SubFeat}_{3it} + \beta_8' \overrightarrow{Reg}_{it} \overrightarrow{Sale}_{it} + \varepsilon_{it}$$

Here, y_{it} is the vehicle's price-to-MSRP, where the subscript i indexes the individual sale records of particular vehicles and t is the month that the sale record takes place. On the right-hand side of the equation, α is a constant term. Of the explanatory variables, \overrightarrow{Feat} is a vector of time-invariant car feature variables including sale region, number of doors, engine liters, number of cylinders, drive type, body type, sale type, fuel type,

induction type, exterior color, interior color, and vehicle subsegment; $\overrightarrow{TimeFeat}$ is a vector of variables that move with time, including age, age², seasonality dummies, and mileage per year; \overrightarrow{Macro} is a vector of macroeconomic variables including the Manheim Index, the U.S. unemployment rate, and an eight-month lag on the automobile inventory-to-sales ratio, as well as year-over-year percent changes in gasoline prices, a 12-month lag of sales of new cars and light trucks, an eight-month lag of sales of new cars and light trucks, and disposable income lagged one month; \overrightarrow{Seg} represents the subsegment of vehicles; $\overrightarrow{FeatMacro}$ is a vector of features and macroeconomic variables interacted with the subsegment of the vehicle, which includes age, MPY, the Manheim Index, the U.S. unemployment rate, and an eight-month lag on the automobile inventory-to-sales ratio, as well as year-over-year percent changes in gasoline price inflation, a 12-month lag of sales of new cars and light trucks, an eight-month lag of sales of new cars and light trucks, and disposable income lagged one month; $\Delta U E$ is the year-over-year change in the unemployment rate; $\overrightarrow{SubFeat}_1$ is a vector of features interacted with the year-over-year change in the unemployment rate, including region of sale and body type; ΔGas represents the year-over-year change in gasoline prices; $\overrightarrow{SubFeat}_2$ is a vector of features interacted with gasoline price inflation, including fuel type, drive type, and region of sale; \overrightarrow{Debt} is the debt service burden in the U.S.; $\overrightarrow{SubFeat}_3$ is a vector of features interacted with the U.S. debt service burden that includes fuel type, sale type and body type; \overrightarrow{Reg} is the growth rate of new-vehicle registrations; \overrightarrow{Sale} represents sale type; and ε_{it} is an assumed i.i.d. Gaussian error term indexed by transaction and time.

The Challenger Model

The Challenger model's methodology can be divided into two separate steps. The first is a ranking step, which provides a continuous percentile ranking to all vehicles of the same make, model and model year sold within a particular month. This first step of the Challenger model provides a quantile forecast for each vehicle based on its characteristics. The

second step runs a group of time series regressions, each of which is indexed to a discrete quantile of price-to-MSRP. The first stage maps VINs to quantiles, and the second stage maps quantiles to price-to-MSRP forecasts. The evolution of a vehicle's price-to-MSRP forecast over time is driven by both its changing quantile in the relevant distribution of cars and the changing price-to-MSRP projection for that quantile over time.

We now describe the two steps of the Challenger model in somewhat greater detail.

The ranking step of the Challenger model uses the generalized linear model to regress the independent variable of percentile rank by make, model, model year and month-of-sale on a host of descriptive vehicle and macroeconomic variables. The percentile rank is calculated by numerically ranking each vehicle 1 through N by residual value-to-MSRP, with N being highest and 1 being the lowest. This number is then converted to an empirical quantile by dividing it by the total number of observations in that particular make, model, model year and sale month, N . This process keeps the auto with the highest residual value as 100%, and all other similarly sorted vehicles fall somewhere in the distribution above zero and below 100%.

The GLM regression used to minimize the errors employs the logit link function for the ranking step. The logit link is appropriate because it constraints quantile forecasts to the interval (0,1). The macroeconomic variables chosen, as in the case of the champion model, attempt to capture drivers of both vehicle demand and supply. A representation of the first stage of the model can be seen in equation (2) below.

Challenger Model: Step 1 Equation

$$(2) \ln(q_{it}/(1 - q_{it})) = \alpha + \beta_1' \overrightarrow{Feat}_{it} + \beta_2' \Delta U E_{it} \overrightarrow{SubFeat}_{1it} + \beta_3' \Delta Gas_{it} \overrightarrow{SubFeat}_{2it} + \beta_4' \overrightarrow{Debt}_{it} \overrightarrow{SubFeat}_{3it} + \beta_5' \overrightarrow{Reg}_{it} \overrightarrow{Sale}_{it} + \varepsilon_{it}$$

Here, q_{it} is the vehicle's empirical quantile, computed as its rank within make, model, model year and month sold divided by the total number of vehicles of that particular

make, model and model year. The subscript i indexes the individual sale records of particular vehicles, and t indexes the month that the sale record took place. On the equation's right-hand side, α is a constant term; \overrightarrow{Feat} is a vector of car feature variables including sale region, number of doors, engine liters, number of cylinders, drive type, body type, sale type, fuel type, induction type, exterior color, interior color, vehicle subsegment, and mileage per year; ΔUE is the year-over-year change in the unemployment rate; $\overrightarrow{SubFeat}_1$ is a vector of features interacted with the year-over-year change in the unemployment rate including region of sale and body type; ΔGas is the year-over-year change in gasoline prices; $\overrightarrow{SubFeat}_2$ is a vector of features interacted with the change in gasoline price inflation including fuel type, drive type, and region of sale; $Debt$ is the U.S. debt service burden; $\overrightarrow{SubFeat}_3$ is a vector of features interacted with debt service burden including fuel type, sale type and body; Reg is the year-on-year growth rate in new-vehicle registrations; \overrightarrow{Sale} represents sale type; and ε_{it} represents a set of normally distributed statistically independent error terms.

The second step of the Challenger model consists of running 11 regressions, one for each decile from 10 to 90 inclusive, as well as two addition regressions for the extreme first and 99th quantiles. Eleven separate percentiles by make, model, model year, and month of sale are created as the independent variables. The prices-to-MSRPs of the 99th, 90th, 80th, 70th, 60th, 50th, 40th, 30th, 20th,

10th and first percentile for each grouping are used on the left-hand side. We then perform an ordinary least squares regression of each price-to-MSRP percentile on a set of macroeconomic drivers. From these regressions we get the behavior of each slice of the price distribution over time, which is then linked to the rankings from the first step to get the residual value for a particular vehicle. The Step 2 regressions for the Challenger model take the form shown in Equation (3) below.

Challenger Model: Step 2 Equation

$$(3) \quad \ln\left(\frac{y_{qjt}}{1 - y_{qjt}}\right) = \alpha + \beta_1' \overrightarrow{TimeFeat}_{qjt} + \beta_2' \overrightarrow{Macro}_{qjt} + \beta_3' \overrightarrow{Seg}_{qjt} \overrightarrow{TimeFeatMacro}_{qjt} + \varepsilon_{qjt}$$

Here, y_{qjt} is the price-to-MSRP by category j , where q indexes the car's quantile, j indexes the category affiliation of the car, where categories comprise a car make-model-year triad; t indexes the month that the sale record takes place; α is a constant term; $\overrightarrow{TimeFeat}$ is a vector of time-varying variables that includes age, age², seasonality dummies, and mileage per year. \overrightarrow{Macro} is a vector of macroeconomic variables that includes the Manheim Index, the U.S. unemployment rate, and an eight-month lag on the automobile inventory-to-sales ratio, as well as year-over-year percent changes in the following: gasoline prices, a 12-month lag of sales of new cars and light trucks, an eight-month lag of sales of new cars and light trucks, and disposable income lagged one month. \overrightarrow{Seg} represents the subsegment of vehicles; $\overrightarrow{TimeFeatMacro}$

represents a vector of time-varying features and macroeconomic variables interacted with subsegments including age, seasonality dummies, mileage per year, the Manheim Index, the U.S. unemployment rate, and an eight-month lag on the automobile inventory-to-sales ratio, as well as year-over-year percent changes in the following: gasoline prices, a 12-month lag of sales of new cars and light trucks, an eight-month lag of sales of new cars and light trucks, and disposable income lagged one month; and ε_{qjt} is a normal, i.i.d. error term.

In practice, to avoid the crossing of forecasts for different quantiles, we estimate Equation (3) for a given category j for $q=0.50$ and then model prices-to-MSRPs for other quantiles of cars in that category by modeling differences between the quantile of interest at quantile $q=0.50$. We interpolate results for cars that fall between the discrete set of estimated quantiles in stage 2.

It should be noted that the Challenger model takes a distinct approach that we would expect might perform better on some aspects of the problem than AutoCycle (e.g. percentile sorting of cars as a function of features). Whether these relative advantages will lead to superior out-of-sample forecasting performance on VIN-level data is fundamentally an empirical question. Based on our prior experience we believe the race will be a very close one. Either AutoCycle or our Challenger model could easily be used as a starting point for modeling and forecasting residual values under alternative economic scenarios.

3. Data

The data we used for model development contain vehicles with model years from 1997 through 2016 inclusive as well as a sample of observed auction sales that have occurred since 2008. The sample contains more than 31 million observations, more than 1,000 vehicle models, and 70-plus car manufacturers. Interested readers can find a list of descriptive statistics of our data in the Appendix. The data were obtained from auction records compiled and provided by the National Automobile Dealers Association.

The NADA dataset contains more than 31 million sale records. To clean the data prior to estimation, we dropped observations for model years before 1997 and outliers with prices-to-MSRPs greater than 120%. This amounts to a 0.3% reduction of the sample. In practice, because our models use the logit transformation of the dependent variable, all vehicles with greater than or equal to 100% of residual value are excluded from the estimation sample. Such vehicles are still included, however, in the forecast sample. These

approximately 130,000 observations account for 0.4% of the remaining sample. The dataset also contains missing observations for a number of variables, but these missing observations occurred because of missing data on car features, rather than missing sale prices. For example, many cars were missing exterior color or induction type but still included the necessary sale price information.

Nearly all sale records in the dataset contain an ample set of time-invariant information about the vehicle, including the vehicle

identification number to the 11th digit. Along with the VIN, the dataset provides information on drive type (four- or two-wheel), fuel type (gas, diesel, etc.), sale type (dealer, salvage, etc.), sale region (Southeast, West, etc.), induction type (turbo, non), and subsegment (luxury, midsize, etc.). Additionally, the dataset provides information on exterior color, interior color and trim. Exterior color was recoded to fit all of the various color names reported in the data into 11 distinct groups, and interior color types were pared down into seven groups. As an example, consider a car listed as the color charcoal—this car would be reclassified as gray. Unfortunately, an exhaustive category reduction scheme proved impossible for trim, because trim can be defined in many different ways by different manufacturers. For example, some manufacturers reference a trim with a sunroof by "EX", while for a different automaker "EX" means trim that includes a spoiler.

Our auction sale data include the week of sale and the mileage of the vehicle at time of sale. From these two pieces of information we can generate other important variables, including average mileage per year and vehicle age at time of sale. We define MPY as the car's total mileage divided by its age. As discussed previously, average MPY, rather than total mileage, is used as a driver in our AutoCycle and Challenger models. The age variable was developed as both yearly age and monthly age using sale week and model year. Since new model year vehicles are often released and sold prior to the start of their official model year, we created age variables that begin in February of the year prior to the model year. For example, the age in months of a 2008 vintage vehicle will be 13 months at the end of February 2008, and the age in years for a model year 2008 vehicle will be one year in February 2008.

The macroeconomic variables and forecasts used in model development and pro-

duction are provided by Moody's Analytics through its DataBuffet platform and sourced from numerous private sector and government institutions. The Appendix includes a table providing the list of macroeconomic variables and their original sources. We forecast car prices under a variety of assumptions, with three macro scenarios being the most important: Baseline, S4 and S6. The Baseline scenario encodes our belief about the most likely path of the economy over the time horizon, the S4 scenario corresponds to a severe recession comparable to the Great Recession, and the S6 scenario revolves around a significant spike in oil prices. Bear in mind that the model can be used to assess any scenario (including all relevant regulatory and client-generated scenarios) so long as all necessary economic variables are populated. In this paper, for ease of exposition, we focus on results for these three scenarios only.

4. Results

In this section, we provide evidence pertaining to four aspects of model performance: core model functionality including aging and mileage effects on car value,

car price forecasts conditional on macroeconomic scenarios, an application of our model to the economics of used-car leasing, and a validation of out-of-sample model

forecasts against forecasts from a strong challenger model.

4.1 Core functionality: Aging and mileage effects

To get a glimpse of the model's core functionalities, we examine aging and usage effects for four cars: a 2013 Toyota Tundra, a 2005 Honda Civic, a 2013 Ford Explorer, and a 2014 Subaru Legacy. All price-to-MSRP forecasts are conditional on the baseline macroeconomic scenario. Age effects are apparent from the steady decrease over time of car values under baseline macroeconomic conditions. We focus on three different usage scenarios, which are as follows:


1. **Stable MPY:** In this usage scenario, mileage increases each month to keep annualized mileage per year stable at its historical average throughout the entire forecast period. The historical average annualized MPY for the entire dataset is 11,526 miles per year.
2. **Increasing MPY:** In the increasing MPY-usage scenario, the car's annualized MPY migrates linearly from the mean to the 99th percentile values in-sample of the mileage-per-year variable over the 48 months of the forecast period. To do this, we took the difference between the annualized 99th percentile of the miles-per-year variable in the dataset—27,030—and the average value of this variable—11,526—and divided that difference by the number of months—48—in the forecast period: $X = (99p\text{-mean})/48$. Car mileage under the scenario is incremented each month by an amount that increases the car's projected usage (measured in miles per year) by X each month.
3. **Decreasing MPY:** In the decreasing MPY scenario, mileage stays fixed at its last observed value, and age in months increases throughout the forecast period. This scenario corresponds to a car that is sitting on the lot without being driven.

Chart 1 displays price-to-MSRP forecasts for the 2013 Toyota Tundra under the increasing mileage-per-year and stable mileage-per-year scenarios. Under the increasing MPY scenario, the Tundra depreciates more quickly. The end result after four years of heavy driving is that the price is 20 percentage points of MSRP lower than it would have been under average usage.

In Chart 2, we display results for a 2005 Honda Civic under the increasing mileage-per-year scenario and the decreasing

Chart 1: AutoCycle Captures Mileage Changes


2013 Toyota Tundra, price-to-MSRP



Sources: NADA, Moody's Analytics

Chart 2: Mileage Effect Less Important With Age


2005 Honda Civic, price-to-MSRP



Sources: NADA, Moody's Analytics

Chart 3: Mileage Overpowers Features


2013 Ford Explorer, turbo vs non-turbo, price-to-MSRP



Sources: NADA, Moody's Analytics

Chart 4: Decreasing Versus Constant Mileage

2014 Subaru Legacy, price-to-MSRP



Sources: NADA, Moody's Analytics

mileage-per-year scenario. The initial price-to-MSRP of the Civic is around 0.13, which reflects the car's relatively old age of 11 years at the beginning of the scenario. The projected price-to-MSRP at the 48-month horizon for the Civic is 0.053 under the decreasing MPY scenario and 0.02 under the increasing MPY scenario. Under both scenarios, but especially the latter, the car is close to being fully depreciated.

Chart 3 displays price-to-MSRP forecasts for the 2013 Ford Explorer turbo under the increasing MPY scenario and forecasts

for the same car but with a non-turbo engine under the decreasing MPY scenario. Observe that the initial price-to-MSRP is 0.05 units higher for the Explorer with the turbo engine compared with its non-turbo counterpart. Over time, however, the price-to-MSRP forecast for the turbo crosses and drops below that of the non-turbo because of the fact that it is being driven much more heavily. The flip in the ordering of the vehicles' prices-to-MSRPs occurs in year one of the forecast, although the reordering of their actual residual values (not shown) oc-

curs later because the Ford Explorer turbo's initial MSRP is higher.

In Chart 4, we display results for a 2014 Subaru Legacy under the stable mileage-per-year scenario and the decreasing mileage-per-year scenario. As in previous charts, we see the effect of no driving versus normal driving leads to an increase in the differential between the price-to-MSRP forecasts of the car under the two scenarios over time. After four years, the Legacy would be 10% of MSRP more valuable if it had been maintained but otherwise sitting on the lot during those four years.

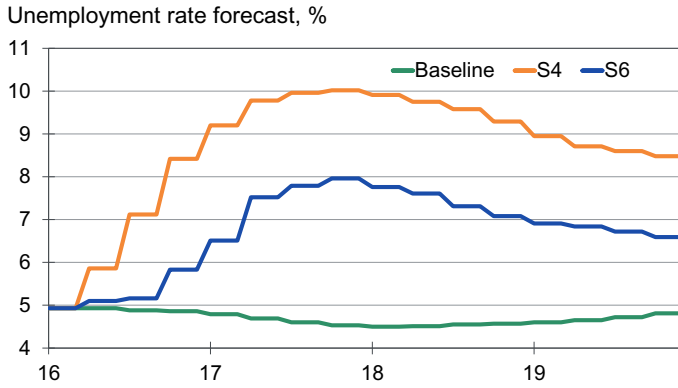
4.2 Forecasting car prices under macroeconomic scenarios

A distinguishing feature of AutoCycle is its ability to forecast car prices under a variety of economic scenarios, including custom macroeconomic scenarios as well as those released periodically by Moody's Analytics and regulators. In the following examples,

we focus on three scenarios: a baseline macroeconomic scenario, a recession scenario (S4), and an oil price shock scenario (S6). The time paths for the unemployment rate and gas prices under each of these scenarios are shown in Charts 5 and 6, respectively.


As shown in Chart 5, unemployment falls slightly and then begins to rise under the baseline scenario. It rises sharply and then recovers under the S4 recession scenario and does essentially the same thing under the S6 oil price shock scenario, but

Chart 5: Unemployment Jumps in S4 and S6



Sources: BLS, Moody's Analytics

Chart 6: Gas Prices Fall in S4 and Jump in S6



Sources: BLS, Moody's Analytics

the onset and magnitude of the increase is less in that case compared with the recession scenario. Gas prices, as shown in Chart 6, rise modestly under the baseline, fall and then recover under the recession scenario, and spike substantially before falling under the oil price shock scenario. Gas prices under S6, it should be noted, fall back below those witnessed under the other two scenarios in early 2018.

The charts depicting car price-to-MSRP forecasts in this section illustrate how heterogeneous, carefully selected pairs of cars can "trade places" under alternative macroeconomic scenarios.

In Chart 7, we show price-to-MSRP forecasts for a 2013 Toyota Corolla versus a 2009 Ford F150. These cars have the same price-to-MSRP of 0.4 at the beginning of the forecast period. Under the baseline scenario, the F150 depreciates slightly faster than the Corolla at first, although the price-to-MSRP

at the end of 2019 is the same for both cars. Under the high oil price scenario, in contrast, the Corolla retains its value far better than the F150 during the first year of the scenario as gas prices skyrocket. This is intuitive: The Corolla has better gas mileage, and its use-cost would be lower than that of the F150 in such a situation. After the point where the gas price under S6 drops below its value under the baseline, the price-to-MSRP forecasts of the two cars cross, as expected, with the price-to-MSRP of the F150 rising briefly as gas prices fall sharply.


Chart 8 depicts price-to-MSRP forecasts for a 2010 Toyota Prius and the same 2009 Ford F150 in Chart 7. Again, the initial price-to-MSRP for both cars is around 0.4, and we compare their paths under the baseline and gas price scenarios. The main difference between the behavior of the Corolla in Chart 7 and the Prius in Chart 8 is that the Prius experiences a more pronounced price increase

than the Corolla under a sharp increase in gas prices and retains slightly more value than the Corolla at a four-year time horizon under the baseline scenario. This makes sense, as the Prius is substantially more fuel-efficient than even the Corolla.

In Chart 9, we show behavior under baseline and high gas price scenarios for a 2010 Toyota Prius and a 2013 Toyota Corolla. This head-to-head comparison makes clear the lessons on the Prius versus the Corolla gleaned from Charts 7 and 8: The Prius retains slightly more value than the Corolla under both scenarios, and the price goes up more in response to a large positive gas price shock.


Chart 10 depicts price-to-MSRP forecasts for the 2014 Chevy Volt versus the 2012 Ford Explorer under the baseline and S6. Although the prices-to-MSRPs of the two cars do not begin at the same value—the Explorer begins at around 0.55 while the Volt starts at 0.4—the subsequent evolution of the car values

Chart 7: Corolla Retains Value During Gas Spike



Sources: NADA, Moody's Analytics


Chart 8: Prius and F150 Part Ways Under S6



Sources: NADA, Moody's Analytics

Chart 9: Prius Trumps Corolla Under S6

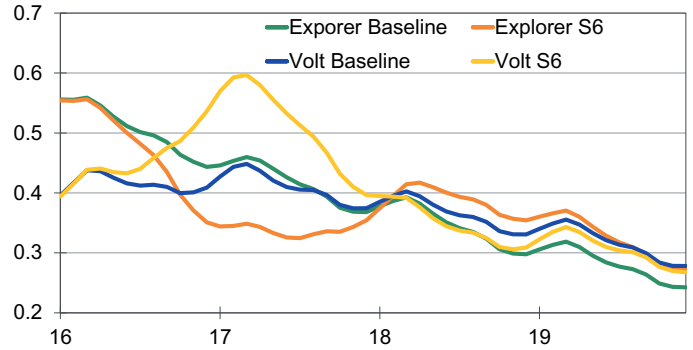
2010 Toyota Prius vs. 2013 Toyota Corolla, price-to-MSRP



Sources: NADA, Moody's Analytics

Chart 10: Volt Charges Past Explorer Under S6


2014 Chevy Volt vs. 2012 Ford Explorer, price-to-MSRP



Sources: NADA, Moody's Analytics

Chart 11: Luxury Cars Hit Hard by Recession


2012 Mercedes SL vs. 2010 Toyota Camry, price-to-MSRP



Sources: NADA, Moody's Analytics

Chart 12: Hybrid Gains More Value Under S6

2012 Toyota Camry, hybrid vs. gas, price-to-MSRP



Sources: NADA, Moody's Analytics

under the S6 scenario has the price-to-MSRP order flipping in late 2016. Under the gas price shock, the Volt quickly becomes more valuable on a price-to-MSRP basis. This differential reverses again in early 2018, after gas prices have fallen. Under the baseline scenario, the price-to-MSRP paths for the two cars cross once, in late 2017, when the Volt overtakes the Explorer.

To change gears slightly, Chart 11 depicts price-to-MSRP forecasts for the 2012 Mercedes SL and the 2010 Toyota Camry under the baseline and recession scenarios. The cars begin at the same price-to-MSRP:

around 0.45. The differential effect of a recession on the price-to-MSRP of each car is distinct, however, with the SL falling substantially in value during the first year of the recession compared with its path during the baseline, while the Camry displays a nearly identical price-to-MSRP path under the two macroeconomic scenarios. This makes sense: Luxury used cars such as the Mercedes often fall in value during recessions as consumers substitute to lower-price cars with a longer expected remaining life such as the Camry.

Finally, in Chart 12 we show the path of a 2012 Toyota Camry Hybrid versus a

2012 Toyota regular gas-powered car under baseline and macroeconomic scenarios. This comparison highlights differences in how the propulsion systems, hence the user cost, can make a difference in car valuation under different economic scenarios, even when the cars are otherwise nearly identical. Two simple conclusions emerge from an inspection of Chart 12: The hybrid retains more value under both scenarios than the gas-powered vehicle, and its price-to-MSRP rises more during the spike in gas prices that occurs under S6 compared with its path under the baseline.

4.3 Accounting for vehicle quality: Forecasting residual values for vehicles by quantiles of quality and unobserved trim

When forecasting a vehicle's residual value, it would be desirable to incorporate information on the vehicle's quality in relation to its peers. Whereas our discussion has centered on conditional expected price-to-MSRP forecasts, we now propose and illustrate a method for conditioning vehicle forecasts on a particular quantile of the distribution of vehicle quality.

In particular, "vehicle quality" here means quality in relation to a group of vehicles of the same make, model and model year as the reference VIN in question. The four steps of the quality adjustment algorithm can be summarized as: (1) MMY portfolio formation, (2) hedonic adjustment of non-reference cars, (3) empirical distribution formation, and (4) marking-to-model of the distribution mean. We will unpack these steps one by one.

In the first step, given the VIN of the reference vehicle, we extract the set of all vehicles in our dataset with the same make, model and model year. We call this set the make-model-year portfolio of the reference car. Second, given the set of observable characteristics of the reference vehicle, we use the AutoCycle hedonic pricing model to adjust the residual value forecasts of all the other vehicles in the MMY portfolio to the VIN-level forecasts that we would obtain if

each of those vehicles had exactly the same set of characteristics as the reference vehicle. In the third step, we extract the model residuals (estimated pricing errors) from the pricing model and add them to the conditional expected price-to-MSRP of the reference car. This third step yields a distribution of observed car prices-to-MSRPs that would have been obtained at auction if all cars in the MMY portfolio had a set of observable characteristics identical to those of the reference vehicle. Fourth, we re-center the mean of the distribution from step three so that it is equal to the conditional expected value of the price-to-MSRP from AutoCycle for the reference car.

Our rich dataset allows us to use the empirical distribution directly to extract values at different quantiles, and our marking-to-model of the distribution mean ensures that our quantile forecasts are consistent with the conditional mean forecasts of our highly accurate hedonic model.

As a caveat to our procedure, we note that non-zero residuals in the AutoCycle model can be ascribed primarily to two sources: variations in used-car vehicle quality and variations in unobserved trim. We believe that the lion's share of variation in the model residuals is the result of quality differences, while trim variations explain a more


modest portion of variations in the pricing error. Users of our quality-dependent price forecasts should proceed with this caveat in mind.

To illustrate AutoCycle's quality-based price-to-MSRP forecasts, let us take as an example the case of a 2009 Ford F150 with a reference VIN equal to 1FTP-W14V89F000000. The MMY portfolio associated with this VIN is the Ford-F150-2009 portfolio. After applying the above algorithm to the cars in this portfolio for the sale month January 2016 we arrive at the histogram of prices-to-MSRPs depicted in Chart 13. Chart 14 depicts the distribution of prices-to-MSRPs for the same MMY portfolio three years later, for the sale month January 2019.

The leftward shift of the distribution during the three-year time horizon separating the distributions in the two charts is evident and occurs because of the (common) aging of the cars in the MMY portfolio. Using the above quality-adjustment algorithm, AutoCycle can compute the price-to-MSRP corresponding to any percentile of vehicle quality. In our used-car leasing case study, which we present in Section 4.5, we will show how adjusting residual value forecasts for quality can impact the minimum profitable lease payment on a used car.

Chart 13: Quality Distribution for January 2016


Y-axis=frequency density, X-axis=% of MSRP, 2009 Ford F150



Sources: NADA, Moody's Analytics

Chart 14: Quality Distribution for January 2019

Y-axis=frequency density, X-axis=% of MSRP, 2009 Ford F150



Sources: NADA, Moody's Analytics

4.4 Forecasting residual values for vehicles of future vintages

Before we proceed to the case study, we first illustrate another key feature of the model: its ability to forecast residual values for future model years. To borrow a term from oenology, let us think of the model year of a car as its vintage. Forecasting residual values for cars of future vintages reduces to the problem of forecasting vintage quality. To forecast vintage quality, we analyze the time series behavior of the residuals ε_{it} from the hedonic pricing model in Equation 1 for each make-model portfolio of cars. This approach allows us to extract the time series properties of the vintage-quality component of prices-to-MSRPs obtained after conditioning on vehicle features and the prevailing macroeconomic conditions when the cars were sold.

Our vintage-quality forecasting algorithm proceeds in three steps. First, for each make-model subset of the MMY portfolios from Section 4.3, we compute the average value of ε_{it} for each separate model year Y . This provides, for each time-invariant MM combination, a time series of estimated average vintage-quality "adjustments" in model year Y . These average residual values are denoted by $\bar{\varepsilon}_Y^{MM}$. In the second step of the algorithm, we compute the vintage time averages of the $\bar{\varepsilon}_Y^{MM}$ residuals for each make-model combination and run an AR(1) model to estimate the speed of reversion of each $\bar{\varepsilon}_Y^{MM}$ series back to its historical mean. We find that on average across MM categories, mean rever-

sion is fairly quick: The average mean reversion coefficient is around 0.49. The third step of the algorithm is projection. In this step, our mean-reverting forecast for vintage quality simply closes the gap between the last estimated value $\bar{\varepsilon}_T^{MM}$ and the MM category time average, $\frac{1}{T} \sum_{t=1}^T \bar{\varepsilon}_t^{MM}$, at the aforementioned rate of mean reversion using an AR(1) model on the de-meaned series.

The intuition behind our vintage quality forecasting algorithm is straightforward. We assume that for the purpose of short- to medium-term forecasts, the average MM quality adjustment is constant relative to other make-models, and that the quality of vintages for a given MM category will revert to this mean reasonably quickly during the next several years. We can apply this vintage-quality overlay to forecast the values of vehicles of future vintages in a way that is likely to have a small impact on overall model accuracy out-of-sample. Furthermore, the overlay approach can be applied to generate residual-value forecasts for vehicles with nearly any combination of previously observed features and any quality level relative to its peers.


Examples of our vintage-quality overlay forecasts are provided in Charts 15 and 16. Chart 15 forecasts the prices-to-MSRPs of a Ford Explorer of four different model years—2017, 2018, 2019 and 2020—during the period 2016-2019. Chart 16 forecasts the

prices-to-MSRPs of the same four vintages of a Toyota Prius during the same period. Both charts depict forecasts obtained under the baseline macroeconomic scenario.

Aging, seasonality and macro effects drive the price-to-MSRP paths of a given vintage of each vehicle in calendar time, as can be seen from the downward sloping price forecasts for each vintage. The initial price-to-MSRP values for each vintage, on the other hand, are determined by a combination of macro effects and the vintage-quality overlay. The baseline macro scenario, which involves a rise in gas prices in 2017, explains the fall of the initial price-to-MSRP for the Explorer in 2017 relative to 2016 as well as the rise in the initial price-to-MSRP for the more fuel-efficient Prius for these same vintages. The role of the vintage-quality overlay is secondary relative to the evolution of macro factors in driving the forecasts. For the Explorer, the vintage-quality adjustment is equal to -0.13 in 2015 and rises over the four-year forecast horizon toward its historical mean of 0.011. For the Prius, the adjustment is -0.04 in 2015 and rises toward 0.00 over the forecast horizon. Thus, the vintage-quality trend for both cars is improving. In the case of the Ford Explorer, however, this small upward trend in vintage quality is more than offset by the macro-driven fall in residual values because of the projected increase in gas prices under the scenario.

Chart 15: Shifts in Future Explorer Vintages


Ford Explorer, model yrs 2017-2020, price-to-MSRP



Sources: NADA, Moody's Analytics

Chart 16: Slight Future Prius Vintage Uptrend

Toyota Prius, model yrs 2017-2020, price-to-MSRP



Sources: NADA, Moody's Analytics

4.5 Case study: Leasing used cars

We now present a case study to demonstrate the utility of the AutoCycle model for residual setting in the context of used-car leasing. For our case study, we used information on the specific cars mentioned in a recent Wall Street Journal article on the subject titled "Hot on the Lot: Leasing a Used Car" (Nagesh & Stoll, 2016). As the article makes clear, used-car inventories are soaring. Two dealers mentioned in the article have opted to start leasing used vehicles in their inventories but face a problem: Lease terms are difficult to calculate correctly. It is unclear how to set residual values in lease contracts for cars that are made available to their second (or subsequent) owner. Fortunately, AutoCycle provides accurate residual value forecasts that can be used to set lease terms in this situation.

The article, and our case study, focuses on two specific vehicles. The first is a 2013 Lexus RX350 sold by a car dealer in the western region of the U.S. with 35,700 miles on the odometer. The second car is a 2014 Jeep Wrangler sold to a car dealer in the Northeast region of the U.S. with unspecified mileage. We will assume for our purposes that the Wrangler has the same mileage as the Lexus (initial mileage for this vehicle is not specified in the article). Let us also suppose that, as in the article, both vehicles are being leased for three years and the mileage allowance for both is 15,000 miles per year.

When forecasting car values in AutoCycle, we must specify additional details about the vehicles. For the Lexus, we assume that the vehicle has a white exterior, a tan interior, and an all-wheel drive powertrain. We assume that the Jeep Wrangler has a black exterior, a black interior, and two doors. The number of doors is not an option in the case of the Lexus (all have four doors), and the drivetrain on the Wrangler must be four-wheel drive. The MSRP of the Lexus is \$40,710, and that of the Jeep Wrangler, given these assumptions, is \$22,395.

Assuming a 100% chance that the car will be returned at the end of the lease, the expected net present value, or NPV, of the lease contract for the lessor is given by:

$$NPV(D, X, r, RV(0), RV(s, T), T) = D - RV(0) + \frac{X}{\left(\frac{r}{12}\right)} \left(1 - \frac{1}{\left(1 + \frac{r}{12}\right)^T}\right) + \frac{RV(s, T)}{\left(1 + \frac{r}{12}\right)^T}$$

Here, D is the initial deposit paid by the lessee, X is the amount of the monthly lease payment, r is the discount factor for the lease payments expressed as an annual percentage rate, $RV(0)$ is the value of the car at the beginning of the lease (time 0), T is the lease term in months, and $RV(s, T)$ is the projected salvage value of the vehicle at the lease's end under scenario s when it is returned to the lessor at time T .

We use AutoCycle to compute $RV(s, T)$ as the conditional expected vehicle value at the T month horizon under macroeconomic scenario s . This is accomplished by forecasting the price-to-MSRP and then multiplying the result by the MSRP stated above for each vehicle. We assume that taxes, transaction costs and insurance costs are zero and that the returned vehicle will be of average quality, given its observed characteristics, projected age, and mileage at the time of return.

In keeping with common practice, we set the discount rate at a level consistent with the lease horizon and the lessee's assumed FICO score. For a 36-month horizon and a FICO score of 635, an appropriate discount rate would be around a 4.75% APR. These are the numbers we will use in our example. We assume that the discount rate r appropriately compensates the lessor for the possibility of nonpayment by the lessee during the course of the contract and that the risk due to the uncertainty of $RV(s, T)$ is diversifiable and does not command a premium.

In this simplified setup, we can learn a number of interesting


things about how to price leases. The first and most obvious lesson is that the expected NPV of the lease contract will be lower under macroeconomic scenarios that depress the salvage value of the returned vehicle. Second, holding other factors constant, there are many combinations of deposit amounts D and lease payments X that deliver the lessor a given NPV for the deal. We devote the rest of the case study to examining different residual value forecasts for our two cars under different macroeconomic scenarios, plotting different combinations of D and X that deliver specific lease NPVs, and examining how deviations of vehicle quality from the average can affect lease terms.

Chart 17 depicts raw car price forecasts, in thousands of U.S. dollars, for the 2014 Jeep Wrangler and the 2013 Lexus RX under the baseline macroeconomic scenario. The initial market value of the Jeep is \$17,110, and the initial value of the Lexus is \$28,604. These values converge over time, with the price gap between the two cars narrowing to around \$4,100 at the three-year horizon and \$2,626 at the four-year horizon.

Chart 18 shows the price forecasts for the two cars under S4, the recession scenario. The gap between the prices under a recession is projected to narrow sharply by the end of 2016, which corresponds to the one-year forecast horizon, because of a significant fall in the secondary market price of the Lexus to around \$16,825. Then, as the worst part of the recession begins to pass, the price gap between the two cars widens again dur-

Chart 17: Lexus Depreciates Faster Than Jeep


Residual value forecast, baseline, \$ ths



Sources: NADA, Moody's Analytics

Chart 18: Severe Recession Hits Lexus Harder


Residual value forecast, S4, \$ ths



Sources: NADA, Moody's Analytics

Chart 19: Gas Price Jump Hits Lexus Harder

Residual value forecast, S6, \$ ths



Sources: NADA, Moody's Analytics

ing the second year of the scenario before remaining roughly constant through years three and four. Overall, the model suggests that the Jeep Wrangler holds its value better than the Lexus during the toughest part of the recession but also recovers less when the recession gives way to recovery.

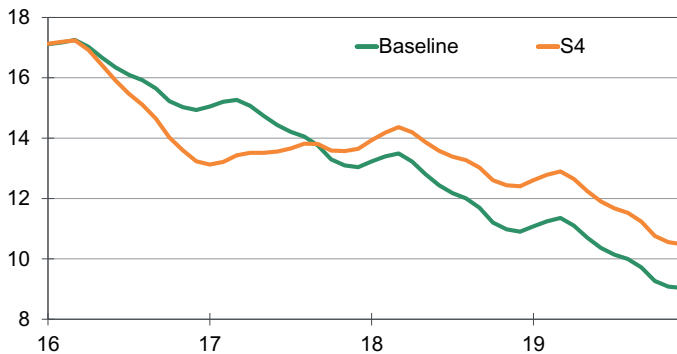
In Chart 19, we show the price dynamics of the two cars under S6, which corresponds to a spike and then a fall in oil and gas prices. In this scenario, the price gap between the two cars narrows to a low point by late in year two as gas prices rise, at which point the prices of both cars recover temporarily as gas

prices fall. The usual aging effects kick in, and the prices of both cars resume their downward trajectory beginning early in year three.

Charts 20 and 21 plot the price paths under the scenarios for the Jeep only, and Charts 22 and 23 plot the price paths under the scenarios for the Lexus only. This allows

Chart 20: Wrangler Has Slight Dip in Recession


Jeep Wrangler, residual value forecast, \$ ths



Sources: NADA, Moody's Analytics

Chart 21: Wrangler Has Shallow Dip From Gas


Jeep Wrangler, residual value forecast, \$ ths



Sources: NADA, Moody's Analytics

Chart 22: Lexus Loses Value During Recession


Lexus RX 350, residual value forecast, \$ ths



Sources: NADA, Moody's Analytics

Chart 23: Lexus Has Significant Dip From Gas

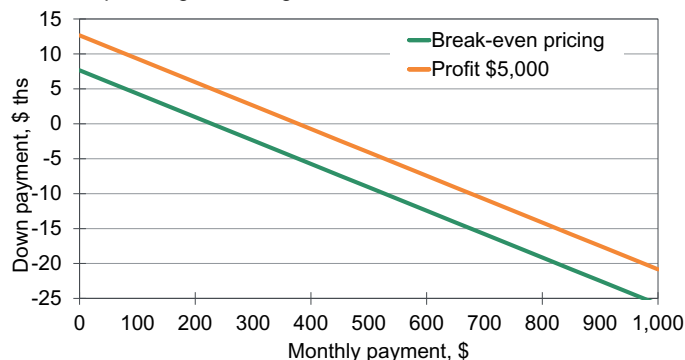
Lexus RX 350, residual value forecast, \$ ths



Sources: NADA, Moody's Analytics

Chart 24: Jeep Lease Pricing


2014 Jeep Wrangler, NPV given 4.75% discount rate



Sources: NADA, Moody's Analytics

Chart 25: Lexus Lease Pricing

2013 Lexus RX 350, NPV given 4.75% discount rate



Sources: NADA, Moody's Analytics

us to discern more clearly the differences between price paths of a single car under alternative scenarios. We see that the Lexus seems quantitatively more sensitive than the Jeep to both a recession and a spike in gas prices, although the directional behavior of both cars is essentially the same.

In Charts 24 and 25, in keeping with the goal of our case study, we plot the different combinations of down payment amounts (*D*) and monthly payment amounts (*X*) that deliver specific net present values to the lessor under the baseline macroeconomic scenario. The down payment is on the y-axis of the chart and the monthly payment amount on the x-axis. The green line corresponds to break-even pricing, or combinations of *D* and *X* for which the net present value of the lease contract to the lessor equals zero. The orange line corresponds to cases in which the NPV to the lessor equals \$5,000. Note that the

orange line lies above the green line in both charts. We can think of this as saying that, for a given monthly payment, if we increase the down payment by \$5,000 and other things are equal, then the contract NPV will rise by \$5,000.

Chart 24 plots the lease pricing loci for the Jeep, and Chart 25 plots the pricing loci for the Lexus. A simple comparison of the charts is instructive. Under a contract with a zero down payment and a target NPV of \$0, the lessor should charge the lessee around \$230 per month to lease the vehicle for 36 months. For the same specifications (*D*=0, *NPV*=0, *T*=36, *s*="Baseline"), on the other hand, the lessor of the Lexus should charge the lessee around \$470 per month, or just more than double the break-even monthly lease payment for the Jeep.

Both calculations assume the same discount rate, of 4.75% APR. Essentially, the

Lexus is more expensive to lease because it loses more dollar value over the lease horizon than the Jeep. Other things being equal, this brings down its contract NPV more, so the Lexus contract requires a higher monthly lease payment to break even. Lessors of either car can make positive profits by increasing the deposit above zero, by charging monthly lease payments above the per-car minimums of \$230 (Jeep) or \$470 (Lexus), or via a combination of these two strategies.

Let us now examine briefly the role of vehicle quality in the determination of break-even lease terms. Table 1 provides price-to-MSRP values corresponding to different deciles of vehicle quality for the 2014 Jeep Wrangler in January 2016 and January 2019. Table 2 provides the same information for the 2013 Lexus RX for deciles 10 to 90 of the quality distribution. We see from comparing the tables that the quality distribution of the

Table 1: 2014 Wrangler Quality Distributions

Jeep-Wrangler-2014	P10	P20	P30	P40	P50	P60	P70	P80	P90
Jan-16	49.80%	57.70%	63.90%	69.50%	74.50%	79.20%	83.30%	87.70%	93.00%
Jan-19	23.10%	29.20%	34.90%	40.80%	46.90%	53.40%	60.10%	68.40%	80.00%

Sources: NADA, Moody's Analytics

Table 2: 2013 Lexus RX Quality Distributions

Lexus-RX-2013	P10	P20	P30	P40	P50	P60	P70	P80	P90
Jan-16	61.00%	63.90%	65.80%	67.10%	68.50%	70.10%	72.10%	74.50%	79.00%
Jan-19	28.30%	30.80%	32.60%	32.90%	34.00%	35.40%	37.10%	39.40%	42.40%

Sources: NADA, Moody's Analytics

Jeep is wider at both time horizons than that of the Lexus.

To obtain the break-even monthly lease payment for each car under different quality assumptions, one must input a quality-adjusted price at time zero and a quality-adjusted price forecast at the end of the lease. For concreteness, consider the case where both cars are at the 60th percentile of their quality distributions at time zero and will drop to the 40th percentiles of the relevant quality distributions after 36 months. According to these assumptions, the Jeep Wrangler's residual value at time zero will be \$17,726, and the Lexus' value at that time will be \$28,527. When the 36-month lease

ends, these values will drop to \$9,137 and \$13,377, respectively. Inputting these numbers into our break-even lease formula under the same specifications as before ($D=0$, $NPV=0$, $T=36$, $s="Baseline"$) yields a break-even lease payment of \$292.61 per month for the Jeep and of \$505.30 for the Lexus. Both payments increase, just as we expect should happen, as the result of the downshift in quality of both cars during the lease term.

Note that, somewhat surprisingly, the 60th percentile value of the Lexus at time zero is marginally lower than the mean value of the Lexus at time zero. This is because the quality distribution of Lexus values is skewed to the right. Nonetheless, the fall

of the Lexus' quality from the 60th to the 40th percentile after 36 months still causes enough additional depreciation in its value to increase the minimum break-even monthly lease payment for the car relative to the case of constant average quality.

Overall, the basic economics of leasing work the same under different economic scenarios and for cars of different qualities. The relative price decreases of different cars, however, will evolve distinctly from the numbers we have used in our examples. Having a lease pricing tool that takes scenario-specific car price forecasts as an input provides a straightforward way of managing residual risk in the face of such possibilities.

4.6 Validation

To validate the AutoCycle model formally, we computed in- and out-of-sample price-to-MSRP forecast performance statistics for the model and a Challenger model. Our in-sample statistics are based on more than 30 million VIN-level observations, and our out-of-sample statistics are based on more than 6.4 million observations during 2015. Our Challenger model was built using best practices for Challenger models, in the sense that we truly did not know a priori whether the AutoCycle or the Challenger model would dominate in out-of-sample performance, and both models were built using the best-known features and techniques pertaining to their model type. In projecting macroeconomic variables for the out-of-sample exercise, we assume perfect foresight of the trajectories of these variables¹ but do not use any of the observations in the out-of-sample period to estimate model coefficients used to form forecasts of prices-to-MSRPs for the out-of-sample cars.

Table 3 displays the in- and out-of-sample R-squared, mean error, mean absolute error, and root mean squared error statistics for the two models, in addition to the

Table 3: AutoCycle 2.0 Versus Challenger


	Model	Obs	R-squared	ME	MAE	RMSE
In-Sample	AutoCycle	30,146,120	0.88	-0.004	0.062	0.083
	Challenger	30,146,120	0.9	-0.001	0.055	0.075
Out-of-Sample	AutoCycle	6,417,497	0.89	-0.006	0.061	0.084
	Challenger	6,417,497	0.87	0.034	0.066	0.089

Sources: NADA, Moody's Analytics

number of observations used to compute them. The Challenger model dominates slightly according to all four of these performance statistics in-sample, whereas the AutoCycle model dominates according to all four statistics out-of-sample for the 12 months of 2015. Furthermore, the objective performance of the model is good: AutoCycle achieves an out-of-sample mean error of only -0.6%, with an out-of-sample R-squared of 0.89. The Challenger model has a larger mean error out-of-sample, of 3.4% (0.034) for comparison, although its R-squared of 0.87 closely lags that of AutoCycle.

Chart 26 plots the out-of-sample RMSE versus time for the AutoCycle model during the 12 months of 2015. The RMSE grows from 8.1% in the first month to 8.6% in the 12th month, or around 0.5% per year. A growing RMSE as a function of the out-of-sample forecasting horizon is typical of forecasting models in many domains.

Chart 26: AutoCycle RMSE Increases With Time



Sources: NADA, Moody's Analytics

¹ Forecasting macroeconomic variables themselves is subject to significant uncertainty, but assuming perfect foresight of such variables in our out-of-sample exercise does not detract from the meaningfulness of our results. Accurate models conditional on such paths simply highlight the importance of considering a reasonable variety of stressed macroeconomic scenarios and understanding the risks they imply for residual values.

In Chart 27, we plot AutoCycle's mean error as a function of the forecast horizon. Around a constant mean of -0.6%, as reported in Table 1, we observe seasonal fluctuations, with observed mean errors lower than average in April and November and around 0% in the

5. Conclusion

In this paper, we introduce a new model for forecasting the residual values of cars at the 11-digit VIN level under a variety of macroeconomic scenarios. We illustrate the plausibility of model forecasts under varying degrees of vehicle usage over multiple time horizons as well as with respect to different macroeconomic scenarios and car features. We then discuss and illustrate two additional features of AutoCycle: (i) the ability to calculate exact car values at a given quantile of inferred car quality (for example, the value of a 2013 Honda Accord at the 80th percentile of quality sold in March 2016, conditional on macroeconomic drivers and the car's observed features); and (ii) the ability to gen-


erate forecasts of residual values for cars in model years that have not yet come on the market (for example, the residual value of a 2019 Honda Accord sold in July 2022).

Following the presentation of the model, our used-car lease case study illustrates a detailed application of our methodology for calculating the break-even minimum monthly lease payments for two cars featured in a recent Wall Street Journal article under baseline projections for a set of macroeconomic variables. We find these break-even amounts, for a Lexus RX 350 and

Jeep Wrangler, to be plausible in light of experience. Finally, we validate our AutoCycle model against a well-performing Challenger model and find that our model is modestly superior to that model—and objectively accurate—out-of-sample given the realized paths of macroeconomic variables in 2015. Our model can be used for residual risk management in large lease and auto portfolios as well as for pricing leases on individual vehicles themselves.

Chart 27: AutoCycle Mean Error Versus Time

Mean error, out-of-sample, in % points



Sources: NADA, Moody's Analytics

Reference

G. Nagesh and J.D. Stoll, "Hot on the Lot: Leasing a Used Car," The Wall Street Journal, April 11, 2016. www.wsj.com/articles/hot-on-the-lot-leasing-a-used-car-1460417861.

Appendix

Descriptive Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
sale_price	30209894	10382.28	8603.267	1	450000
mileage	30209893	81339.58	55430.7	0	999999
model_year	30209894	2006.294	4.228712	1997	2016
msrp	30160687	27217.41	11329.52	0	474600
region	30209894	2.338566	1.066871	1	4
year	30209894	2011.815	2.617909	2007	2015
mage	30209894	83.50119	48.65217	2	239
yage	30209894	6.520649	4.056077	0	19
mileage_per_year	30209893	11623.45	5075.52	0	247604.4
residual_value	30152036	0.374651	0.237284	0.002712	2.903862
Doors	30209894	3.309972	1.241033	1	5
Liters	30209894	7.292512	2.516004	1	16
cylinders	30209894	7.435309	1.328566	1	9
drive_type	30196053	3.020077	1.161569	1	5
body_type	30196053	23.34537	6.626794	2	29
Model	30209894	595.5382	285.0881	1	1142
sale_type	30209894	2.73054	1.00283	1	7
vehicle_type	30196053	1.946444	1.005864	1	4
fuel_type	30196053	3.970475	0.40935	1	7
transmission	30196053	2.8893	0.449029	1	3
induction type	30196053	2.11772	0.480161	1	4
exterior	30209894	2.800391	3.131003	1	11
interior	30209894	1.613232	1.259621	1	7
Make	30209894	29.75958	19.50401	1	72

Macroeconomic Variables and Their Sources

Alias	Description	Sources
FREG	New Vehicle Registrations: Total	R.L. Polk & Co.; Moody's Analytics Forecast
FSDEBT	Debt Service Burden: Total	U.S. Board of Governors of the Federal Reserve System; Moody's Analytics Forecast
FCPIUETB01	CPI: Urban Consumer - Gasoline - All Types	U.S. Bureau of Labor Statistics; Moody's Analytics Forecast
FYPDPIQ	Income: Disposable Personal	U.S. Bureau of Economic Analysis; Moody's Analytics Forecast
FSCARD	Automobile Inventory-to-Sales Ratio	U.S. Bureau of Economic Analysis; Moody's Analytics Forecast
FRVEHL	New Vehicle Sales: Cars and Light Truck	U.S. Bureau of Economic Analysis; Moody's Analytics Forecast
FLBR	Household Survey: Unemployment Rate	U.S. Bureau of Labor Statistics: Current Population Survey; Moody's Analytics Forecast
FMUVIM	Manheim Used Vehicle Value Index	Manheim Consulting; Moody's Analytics Forecast

About the Authors

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, both 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.

Samuel W. Malone is a director in the Specialized Modeling Group at Moody's Analytics. Dr. Malone has taught and consulted at top institutions in Europe and the Americas including Oxford, the University of Navarra, the European Commission, the Central Banks of Venezuela and Peru, and several large North American financial institutions. He is coauthor of the book Macrofinancial Risk Analysis, published in the Wiley Finance series with foreword by Nobel Laureate Robert Merton, as well as the author of numerous academic journal articles in economics and statistics. His articles have been published in outlets such as World Development, the Journal of Applied Econometrics, the Journal of Financial Econometrics, the International Journal of Forecasting, and the Annual Review of Financial Economics. He holds undergraduate degrees in mathematics and economics from Duke University, where he studied as an A.B. Duke scholar and graduated with summa cum laude Latin honors, and master's and PhD degrees in economics from the University of Oxford, where he studied as a Rhodes scholar.

Michael Brisson is an economist at Moody's Analytics. He develops state and local revenue forecasts on various consulting projects and is a frequent contributor on energy-related issues. Mike holds a PhD in Applied Economics from Northeastern University, an MS in Economics from the University of Buffalo, and a BA in Political Science from the State University of New York at Oswego.

Michael Vogan is an economist in the Credit Analytics department at Moody's Analytics. Before joining Moody's Analytics, Michael was a research analyst at the Federal Reserve Bank of Philadelphia. He holds a master's degree in applied economics and econometrics from the University of Delaware and a bachelor's degree in economics from Bloomsburg University.

About Moody's Analytics

Moody's Analytics helps capital markets and credit risk management professionals worldwide respond to an evolving marketplace with confidence. With its team of economists, the company offers unique tools and best practices for measuring and managing risk through expertise and experience in credit analysis, economic research, and financial risk management. By offering leading-edge software and advisory services, as well as the proprietary credit research produced by Moody's Investors Service, Moody's Analytics integrates and customizes its offerings to address specific business challenges.

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.moodyanalytics.com.

About Moody's Corporation

Moody's is an essential component of the global capital markets, providing credit ratings, research, tools and analysis that contribute to transparent and integrated financial markets. **Moody's Corporation** (NYSE: MCO) is the parent company of Moody's Investors Service, which provides credit ratings and research covering debt instruments and securities, and **Moody's Analytics**, which encompasses the growing array of Moody's nonratings businesses, including risk management software for financial institutions, quantitative credit analysis tools, economic research and data services, data and analytical tools for the structured finance market, and training and other professional services. The corporation, which reported revenue of \$3.5 billion in 2015, employs approximately 10,400 people worldwide and maintains a presence in 36 countries.

© 2016, Moody's Analytics, Moody's, and all other names, logos, and icons identifying Moody's Analytics and/or its products and services are trademarks of Moody's Analytics, Inc. or its affiliates. Third-party trademarks referenced herein are the property of their respective owners. All rights reserved. ALL INFORMATION CONTAINED HEREIN IS PROTECTED BY COPYRIGHT LAW AND NONE OF SUCH INFORMATION MAY BE COPIED OR OTHERWISE REPRODUCED, REPACKAGED, FURTHER TRANSMITTED, TRANSFERRED, DISSEMINATED, REDISTRIBUTED OR RESOLD, OR STORED FOR SUBSEQUENT USE FOR ANY PURPOSE, IN WHOLE OR IN PART, IN ANY FORM OR MANNER OR BY ANY MEANS WHATSOEVER, BY ANY PERSON WITHOUT MOODY'S PRIOR WRITTEN CONSENT. All information contained herein is obtained by Moody's from sources believed by it to be accurate and reliable. Because of the possibility of human and mechanical error as well as other factors, however, all information contained herein is provided "AS IS" without warranty of any kind. Under no circumstances shall Moody's have any liability to any person or entity for (a) any loss or damage in whole or in part caused by, resulting from, or relating to, any error (negligent or otherwise) or other circumstance or contingency within or outside the control of Moody's or any of its directors, officers, employees or agents in connection with the procurement, collection, compilation, analysis, interpretation, communication, publication or delivery of any such information, or (b) any direct, indirect, special, consequential, compensatory or incidental damages whatsoever (including without limitation, lost profits), even if Moody's is advised in advance of the possibility of such damages, resulting from the use of or inability to use, any such information. The financial reporting, analysis, projections, observations, and other information contained herein are, and must be construed solely as, statements of opinion and not statements of fact or recommendations to purchase, sell, or hold any securities. NO WARRANTY, EXPRESS OR IMPLIED, AS TO THE ACCURACY, TIMELINESS, COMPLETENESS, MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OF ANY SUCH OPINION OR INFORMATION IS GIVEN OR MADE BY MOODY'S IN ANY FORM OR MANNER WHATSOEVER. Each opinion must be weighed solely as one factor in any investment decision made by or on behalf of any user of the information contained herein, and each such user must accordingly make its own study and evaluation prior to investing.