Off With Their Assists: An Empirical Study of Driving Skill in Forza Motorsport 4

Thomas Debeauvais
Department of Informatics
University of California, Irvine
tdebeauv@uci.edu

Thomas Zimmermann
Nachiappan Nagappan
Microsoft Research, WA
{tzimmer, nachin}@microsoft.com

Kevin Carter, Ryan Cooper,
Dan Greenawalt, Tyson Solberg
Turn 10 Studios, WA
{v-kecart, ryanc, dangre, tysonsol}@microsoft.com

ABSTRACT
This paper describes patterns of skill progression in Forza Motorsport 4, an Xbox 360 racing game. Using in-game telemetry data from more than 200,000 players and 24 million races, we characterize how players use and customize driving assists such as the trajectory line, automatic gear shifting, or assisted braking over time. We find that some of the assists are never disabled by significant player segments. Some "yoyo" players repeatedly enable and disable assists. We also present a model to predict when a player is ready to successfully disable an assist with a precision ranging from 60 to 90 percent.

Author Keywords
assists; difficulty; progression; racing game; Forza

ACM Classification Keywords
K.8.0 Personal Computing: General—*Games*

INTRODUCTION
Skill level and player progression are two integral parts of any game. If the game is too difficult, players are frustrated. If it is too easy, they are bored. Balancing the two ensures players enjoy the overall experience. This is one of the main reasons why game designers often provide an option that lets the players specify the level of difficulty themselves. But players may not always know what level of difficulty will work for them, or they may not be confident enough in their skill to increase the difficulty by themselves.

In this paper we aim to answer several questions about skill level and player progression. We look at the patterns that player progression follows. We also investigate the ability to predict when players are ready to increase the difficulty. Answering these questions is an important first step to improve game design, and also towards recommender and AI systems assisting players to challenge themselves and have more fun.

More specifically, we study in this paper how players of Forza Motorsport 4 (FM4) improve in skill and change the game's difficulty. FM4 lets players customize the game's difficulty by providing several assists. Assists are game mechanisms such as the trajectory line or automatic gear shifting that help a player race. Not all players have the same skill, so players can enable or disable assists to fit their needs with some level of granularity. While disabling some assists may not always increase the game's difficulty, disabling others can make it impossible to drive at all for a player. We look at how these assists are used, and how they are disabled over time, for more than 200,000 players. This allows us to determine which assists are useful to beginners, completely unforgiving, or always-on assists that are rarely disabled. By looking at the trends in players' racing data, we can model which factors determine when players can safely disable an assist without being bored or frustrated.

We first provide a review of the literature on skill, progression, and racing games. To familiarize the reader with FM4, we provide a quick overview of the game with a description of each of the assists. We then present a summary of our data. We explore in depth how players use and disable assists, and propose two models predicting if a player is ready to disable an assist. We conclude with the implications of our findings for the design of racing games and multi-user game-based learning environments.

RELATED WORK
On the qualitative side, Reeves et al. broke down Counter-Strike actions to show that experts are not concerned by particular actions, but rather by chains of actions [1]. Nardi et al. analyzed World of Warcraft chat logs to see how players learn about strategies, item usage, and game ethos through the in-game chat [2]. On the quantitative side, Huang studied the evolution of skill in Halo: Reach players, with a particular focus on regaining skill after breaks [3]. Through standardized Xbox 360 achievements, Phillips compared progression between several shooter games [4]. Kraaijenbrink et al. and Bateman et al. tried adding handicaps to balance games and make them more fun for both experts and less-skilled players alike [5] [6]. More towards racing game development, Hullett et al. found that players use few of the available tracks, cars, and game modes [7]. Finally, Togelius et al. procedurally generate

race tracks based on player driving styles, and propose several definitions of “fun” tracks [8]. Telemetry data has been used for many other purposes to maximize player value [9].

FORZA MOTORSPORT 4

Forza Motorsport 4 (FM4) is an Xbox 360 racing game developed by Turn 10 Studios. FM4 is the fourth installment in the series of Forza games. It was released in October 2011 and has sold four million copies to date [10].

Cars and modes

FM4 is a complex racing simulation where each car has a particular weight, steering rate, drive type (e.g., all-wheel drive, AWD, versus rear-wheel drive, RWD), and so on. To allow players to compare cars at a glance, the game summarizes the performance of a car in a score called the Performance Index (PI) ranging from 0 for small city cars to 1 for purpose-built racing cars. Players can acquire cars by spending in-game credits or as a level-up reward. These credits and experience points are awarded when a player completes a race. Races can take place in several modes:

- In *career mode*, players progress through ten divisions, each containing a dozen races. Higher division races involve cars with higher PI controlled by more difficult AI.
- In *quick race*, players can select any track and car they want. Quick races are offline and involve up to two players sharing the same screen.
- In *online mode*, players can participate in multiplayer activities ranging from the usual lap races pitting up to 16 players to drag races or even car soccer.
- In *rivals*, a single player races against the time that another (remote) player previously set for the race.

Assists

Assists are in-game mechanisms that help the player to drive. The FM4 telemetry logs six assists, each of which can take two or three levels. A seventh assist, steering, makes turning easier, but was not recorded. Turn 10 Studios designed most of the assists to make players more consistent and competitive [11].

Stability control prevents the car from spinning when cornering too fast. Traction control prevents it from spinning when accelerating. When the game detects that the car starts spinning, these two assists will slow the car down. Each of these two assists can be turned *on* or *off*.

The braking assist supports the player when he/she brakes or should brake. It can take three levels: in *assisted with Anti-lock Braking System* (ABS), it automatically slows the car down when approaching a turn. In *non-assisted with ABS*, it prevents the car from drifting when braking in a straight line, but may increase the braking distance. In *non-assisted without ABS*, the player has complete control of the brakes.

<table>
<tr><th></th><th>Easy</th><th>Medium</th><th>Hard</th><th>Advanced</th><th>Expert</th></tr>
<tr><th>Stability</th><td colspan="2">ON</td><td colspan="3">OFF</td></tr>
<tr><th>Traction</th><td colspan="3">ON</td><td colspan="2">OFF</td></tr>
<tr><th>Braking</th><td colspan="2">Assisted w/ ABS</td><td>ABS</td><td colspan="2">OFF</td></tr>
<tr><th>Shifting</th><td colspan="3">Automatic w/o clutch</td><td>Manual w/o clutch</td><td>Manual w/ clutch</td></tr>
<tr><th>Line</th><td colspan="2">Full</td><td colspan="2">Brake</td><td>OFF</td></tr>
<tr><th>Damage</th><td>Cosmetic</td><td colspan="2">Limited</td><td colspan="2">Simulation</td></tr>
</table>

Table 1: Values of assists in the built-in bundles.

The shifting assist helps the player in passing gears. It can take three levels. Gears can be shifted *automatically*, much like in real-life automatic cars. Gears can also be shifted *manually*, in which case the player has to press a button to shift gear up, and another button to shift gear down. Manual shifting gives the player control over the gear ratios, which, if used properly, can result in faster acceleration. Gears can also be switched *manually with clutch*, in which case the player must press the clutch trigger to switch gears.

The line assist overlays the optimal trajectory to follow on the track. It also takes three levels. In *full*, the line shines green when the player should accelerate and red when the player should brake. In *brake*, only the red portions are displayed. The line can also be completely turned *off*.

The damage assist determines how much the performance of the car can change during the race. In *cosmetic*, collisions only leave visible traces on the car. In *limited*, collisions can reduce the car’s performance. In *simulation*, collisions reduce the performance of the car a lot more, tires wear off, and the player has to think about refueling.

We call *bundle* a configuration of assists. The game ships with five built-in bundles: easy (selected by default), medium, hard, advanced, and expert. Table 1 describes the value of assists in the built-in bundles. A custom bundle is a configuration of assists that is not built-in but player-defined. Players can change the bundle or the value of a particular assist in a menu before any race. In career and online modes, each disabled assist increases the credit rewards by 10%, giving players an incentive for disabling an assist.

We provide Figure 1 and a short scenario to give an idea of the skill necessary to drive using the easy and expert bundles: When approaching a turn in *easy*, the player stops accelerating and steers so as to follow the full line. The player “only” coordinates two controls together: the right trigger to accelerate and the left stick to steer. In *expert*, the player guesses the optimal trajectory, shifts gears down in sync with the clutch, and taps the brakes lightly so as not to skid. This involves five simultaneous controls and much more anticipation of the game overall.

Figure 1: Xbox 360 controller and the FM4 controls.

DATA AND DISTRIBUTIONS

Data

Our main dataset consists of the race entries from a random sample of 5% of the whole player base[1]. The 220,000 players generated 24.5 million race entries from the FM4 launch in October 2011 until July 2013. A race entry is created when the player leaves a race, whether abandoned or finished. Each race entry contains the value of the aforementioned assists, and other metrics such as the rank at the finish line and the car's PI. The game telemetry logs each of the 3-value assists into two binary assists (on or off). For example, the 3-value braking assist translates to two binary assists called ABS and autobrake, which are partially dependent: the players who use autobrake also use ABS, but those using ABS do not necessarily use autobrake. Therefore the six assists described in the previous section are represented with ten *binary* assists for the rest of the paper: stability, traction, autobrake, ABS, autoshift, clutch, full line, brake line, cosmetic damage, and limited damage. This abstraction reduces the complexity of the analysis, as now all assists are binary. The telemetry tracked neither the input device (Xbox 360 controller vs steering wheel) nor the view during the race (cockpit vs wheel level vs above car).

To complement this racing data and get a sense of player progression through the game as a whole, we also look at the Xbox Live achievements that players gain over time. Given the large size and representativeness of the dataset, all effects will be significant. Thus we focus on the size of the effects rather than on their significance.

General trends and distributions

An achievement is unlocked every ten driver levels and at the end of each division. Figure 2 plots the number of players who unlocked each achievement, and the average time it took. Half of the players reach driver level 20, where moderate-performance cars are offered. Around 15% reach level 50, where the fastest cars are offered. More than 20% import data from FM3, the previous game in the franchise, and they do so, on average, a few days after they start playing.

Around 48% of all races in our dataset happen in career mode, 30% online, 16% in quick race, and 6% in rivals. Given this distribution, we place more focus in the rest of the paper on the career and online modes than on the other modes. Not all players race equally. The median player has raced 29 times, and the top 5% more than 434 times. These top 5% players account for half of all races. Players race three times per day on average (median 1.2).

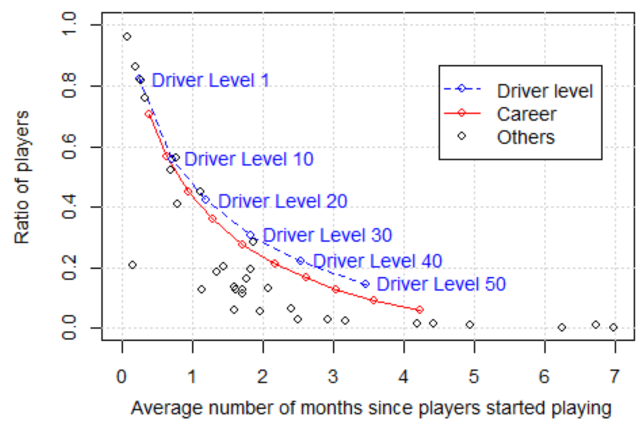

Figure 2: Tracking progression through achievements in FM4.

PATTERNS OF ASSISTS

First-race bundle

The game ships with the easy bundle selected by default. Thus it is not surprising that 85% of players race for the first time using the easy bundle. This also means that 15% of players race for the first time using a bundle different

	Career		Online	
Bundle	%players	%races	%players	%races
Easy	57.9	18.5	29.9	3.0
Stab,trac,clutch,cosm dmg	28.2	1.8	49.5	5.0
Stability, traction, clutch	21.0	2.1	12.5	0.5
Stab, trac, clutch, ltd dmg	13.7	0.8	4.0	1.1
Medium	10.5	3.2	0.8	0.0
All but autobrake	7.6	4.0	12.5	2.8
All but autob and fullline	4.2	2.7	10.0	2.6
Clutch and cosmetic dmg	1.2	0.5	31.8	20.4

Table 2: Bundle popularity across modes. The "%players" column contains the percentage of players who *ever* used the bundle among the players who ever played the given mode. A player can contribute to multiple bundles. The "%races" column contains the percentage of that mode's races with that bundle. A race contributes to only one bundle. (Legend: Stab: Stability; trac:Traction; cosm:Cosmetic; dmg: Damage; ltd:Limited; autob:Autobrake)

---

[1] To avoid sampling bias, we picked players which id modulo 20 was 0. The id is an auto-incremented field generated by the database.

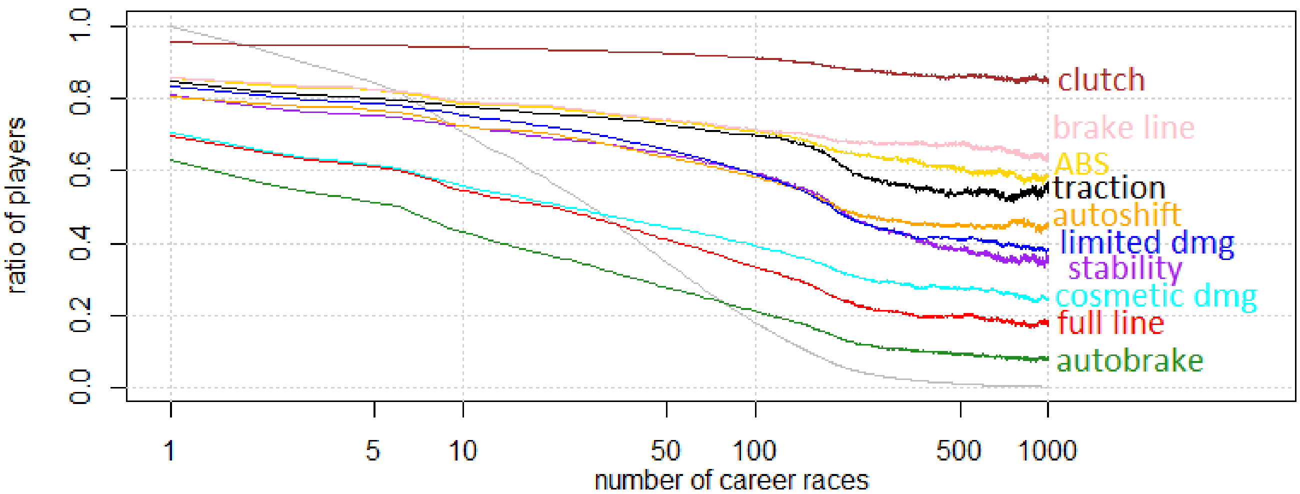

Figure 3: Evolution of each assist over *career* races. The gray line indicates the ratio of players contributing to X career races.

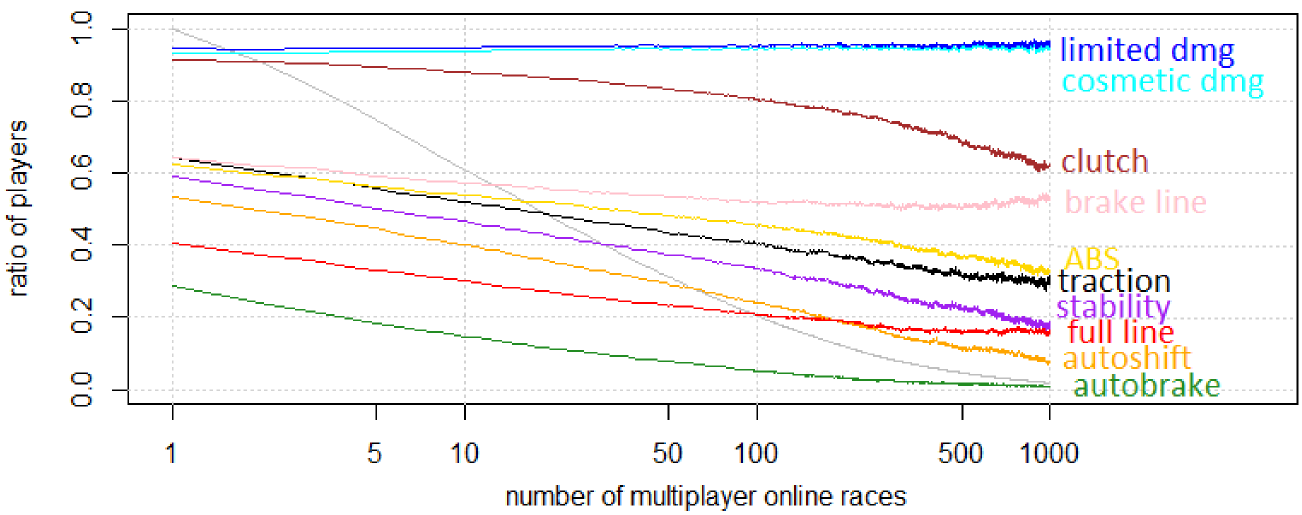

Figure 4: Evolution of each assist over *online* races. The gray line indicates the ratio of players contributing to X online races.

than easy (medium, hard, advanced, expert, as well as custom). Note that the first race, albeit full-fledged with AI opponents and moderate-performance cars, is a one-lap hands-on tutorial in quick race mode, aimed at explaining the controls to the player. Thus it is understandable that all the assists are enabled by default.

Among the players who start with the easy bundle, most eventually change the assists, but a small proportion of players never touch them. The players who eventually change the assists race on average 147 times (median 57) over 163 days (median 81). On average, they change an assist for the first time after 9 races (median 3). Most of the players who never touch the assists race less than 10 times. We call this last segment of players the *"samplers"*, since they only experienced a sample of the game.

Most frequently used bundles

Before looking at individual assists, we look at the bundles commonly used by players. With ten binary assists, there are around a thousand possible bundles. However, only the eight reported in Table 2 are used at least once by 10% of players of any mode. In terms of players, the Easy bundle is by far the most used bundle in career mode, and the third most used in online mode. The large segment of "samplers" mentioned before helps explain this. Among the other built-in bundles, only Medium is somewhat popular in career mode, but most players seem to customize the difficulty at the assist level rather than at the bundle level. Moreover, the popular bundles in online mode keep damage at the cosmetic level, probably because players know they will collide with each other often and do not want their performance to suffer from it. And finally, clutch is the only assist that is enabled in all of the bundles in Table 2.

The player-defined bundle with stability, traction, clutch, and cosmetic damage assists enabled is used by almost a third of career players (28.5%) and half of online players (49.5%). Yet no built-in bundle resembles it.

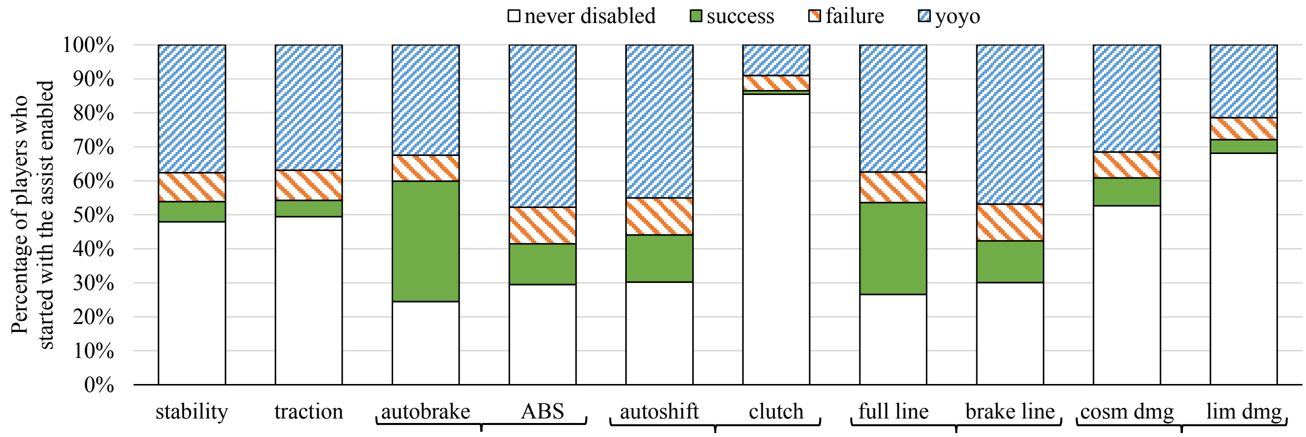

Figure 6: Success, failure, and yoyo for each assist and for the players who ran their first race (in any mode) with the assist enabled.

Assist progression

How many races does the average player need to feel confident about disabling an assist? Once again, we distinguish between the career and online modes. We plot the evolution of assist usage over the number of races in Figure 3 and Figure 4. The graph is cut at 1,000 races because less than 1% of players ever raced more than that in either mode.

In their first career mode race, around 70% of players enable the full line, autobrake, and cosmetic damage, 95% enable the clutch, and 80-85% the other assists. Over time, autobrake and full line are disabled twice faster than other assists, suggesting they are the easiest to disable, or the least enjoyable. Clutch stays flat, suggesting it is the hardest to disable. The slight dip around 200 career races marks the final race of career mode; beyond that point, players replay career races they already completed.

Assists follow the same trends in online mode with several exceptions. First, players start their first online race with more assists disabled than in their first career race. For example, autobrake is used in only 30% of the first online race versus 60% of the first career race. Second, clutch is disabled faster and by more players in online mode. Players who race online may have practiced in career mode beforehand, and seem to be more skilled than career mode players. Yet after a thousand online races, 20% of online players still use the full line, and 50% the brake line. Figure 4 also confirms that players do not want to be penalized by colliding with others: damages stay flat at 95%.

Assist transitions

Once a player disables an assist, several scenarios can happen. (1) The player can keep it disabled forever, in which case the increased difficulty matches the player's skill, and we call the disabling a "success". (2) Otherwise, the player eventually re-enables the assist. If the player never disables the assist again, racing without the assist was so hard that the player never dares doing it again. In this case, we call the disabling of the assist a "failure". (3) If the player re-enables and re-disables the assist once or more, we call the disabling a "yoyo". Figure 5 summarizes these terms and highlights the race before and the race after the first time an assist is disabled.

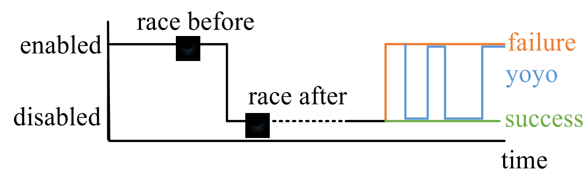

Figure 5: Success, failure, and yoyo after disabling an assist.

For each assist, Figure 6 provides the percentages of successes, failures, yoyos, and players who never disable. Autobrake is easiest to disable with 35% of success, while clutch the most difficult with less than 1% of success. Overall, players who "fail" re-enable an assist after very few races. For example, 38% of the stability failures re-enable it the very next race. This percentage rises to 56% for clutch, 69% for autobrake, and 80% for brake line. Disabling an assist seems to be an immediate hit-or-miss.

MODELING ASSIST TRANSITIONS

To model the success or failure of players in disabling a particular assist, we restrict our sample to the players in the success and failure categories for that assist. We investigate two possible scenarios:

Scenario 1: A player finished a race. Can we *predict* if the player is ready to successfully disable an assist?

Scenario 2: A player finished a race, disabled an assist, and then finished a second race. Can we *characterize* what makes a successful disabling of the assist?

We have one dataset per assist. Each dataset holds racing data from the races before and after the assist was disabled. The features we select aim at measuring in-game racing

Assist	Num races	Races/day	Career Mode	RWD	Car PI	Position
Stability	–	–		+	–	–
Traction	–	–		+	–	
Autobrake	–	–	–	+	–	–
ABS	–	–	–	+	–	–
Autoshift	–	–	–	+	–	–
Clutch	–	–		+	–	
Full line	–	–	–	+	–	–
Brake line	–	–	–	+	–	–
Cosmetic Damage	–	–		+	–	–
Limited Damage	–	–		+	–	–

Table 3: Sign of the log-odds-ratios from a logistic regression for Scenario 1: "race before". A + indicates that the feature is positively correlated with successfully disabling the assist. Empty cells have p>.001.

Assist	Num races	Races/day	Days diff	Career Mode	RWD to FAWD	Car PI	Car PI diff	Finished	Position
Stability	–	–	+	+	+			+	–
Traction	–	–	+	+	+				–
Autobrake	–	–			+			+	–
ABS	–	–	+		+			+	–
Autoshift	–	–	+		+		–	+	–
Clutch	–	–	+	+	+				–
Full line	–	–	+		+		–	+	–
Brake line	–	–	+	+	+			+	–
CosmDmg	–	–	+	+	+	–	+	+	–
LimDmg	–	–	+	+	+	–		+	–

Table 4: Sign of the log-odds-ratios from a logistic regression for Scenario 2: "race before and after". Same legend as Table 3.

skill. They are drawn from our experience playing the game and from the literature on game skill.

For each assist and scenario, we run a multivariate logistic regression using the player's race statistics (see below) as independent variables, and the success of disabling an assist as the dependent variable (outcome). For each independent variable, the regression computes the log of the odds ratio and a significance level. For example, a negative log odds ratio for the number of races means the odds of successfully disabling an assist are higher for players with fewer races, fixing all other independent variables. In other words, the sooner players disable an assist the more likely to succeed.

Once a model is built, it contains one coefficient (the log odds ratio) per variable. We can then input the data of a player who was not used to build the model to predict whether she will disable an assist successfully or not. The models that are built for Scenario 1 can be used as input for recommender systems suggesting which assists to disable; models from Scenario 2 can shed light on the factors contributing to successfully disabling an assist.

Scenario 1: Predicting the success of disabling

To assess if a player were to successfully disable an assist after the completion of a race, we measure player skill using metrics generalizable to any racing game. We use metrics only from the race *before* the assist was disabled.

We saw that the more races played, the more likely to disable an assist; the number of races is our first metric **(Num races)**. The position at the finish line is common to all racing games, and finishing last is a strong signal of underperformance; this is our second metric **(Position)**. When we played the game, we found career mode more appropriate for training than online mode: it is only in career mode that the player can lower the AI's difficulty or rewind the race if she misses a turn **(Career Mode)**. A previous work noted large differences between the number of games and the number of days one has been playing for [3], so we add the number of races played per day **(Races/day)**. We also add two car metrics: the car performance index **(Car PI)** and whether the car's drive type is Rear-Wheel Drive **(RWD)**. Cars with higher PI are more "twitchy" and demanding in terms of reflexes and skill. An FM4 developer admits using the traction assist when driving an RWD car [12]. Cars with high PI numbers are generally RWD.

Table 3 describes the influence of each factor on the model. We expect the players who race only a few times to be less skilled than players who race many times, and thus less likely to keep an assist disabled – they disable too early. However we find the opposite in the data: the odds of successfully disabling an assist are higher for the players who participate in fewer races. In other words, the players who disable an assist early are more likely to keep it disabled. We saw earlier that 20% of the player base played FM3 before. These players already know which assists they need. On the other hand, players who race for a long time with an assist enabled may become used to it.

Scenario 2: Characterizing a successful disabling

We saw earlier that most of the players who re-enable an assist do so the race immediately after they disabled it. In order to assess the factors in the race that influence the re-enabling of an assist we start with factors from model 1: career mode, position at the finish line, number of races so far, and number of races per day so far. Previous literature shows that skill degrades quickly over time [3], so we also add the number of days elapsed between the race before and the race after disabling the assist **(Days diff)**. RWD was also included in the early stages of Model 2, but became insignificant when we added to the model whether players switch from RWD to FAWD **(RWD to FAWD)**. Switching from FAWD to RWD is also insignificant for all assists, and thus not reported. The car performance index **(Car PI)** is carried over from model 1, complemented with the difference in car PI between the race before and the race after **(Car PI diff)**. A positive PI difference indicates that the player races a car that requires more skill than the car used in the previous race. And finally, we add to the model whether the player finishes (as opposed to abandons) the race **(Finished)**.

Controlling for all other factors in the model, the odds of maintaining the assist disabled are higher for players who finish the race. This is not surprising: a player who has no fun without the assist can (and probably should) abandon the race and re-enable it.

In Scenario 1 the odds of successfully disabling an assist are lower for the players who race in career mode *before* disabling the assist. Here in Scenario 2, the odds of success are higher for players who race in career mode *after* disabling the assist. In other words, the players who are the most likely to successfully disable an assist are those who race in online mode, disable the assist, and then race in career mode. Career mode is indeed better for practice.

Evaluating the models

In order to assess the efficacy of our approach at predicting whether a player will keep an assist disabled, we train models for each scenario and each assist on two thirds of the data (picked randomly), and test on the last third. For each test instance, the regression outputs a likelihood between 0 and 1. The cutoff point is set at 0.5. We repeat this process 50 times, and report the average precision and recall for each assist and each model in Table 5 below.

In our case, precision matters more than recall. A low precision means that the model overestimates the skill of a lot of players. If the system tells the player that she is ready, when she actually is not, disabling the assist may lead to an unpleasant game experience. On the other hand, a low recall means that the model underestimates the skill of a lot of players, which is acceptable: if the game is too easy, players can disable the assists themselves, as is presently the case.

Overall the precision is between 0.6 and 0.9 for Scenario 1, suggesting that in at least 6 out of 10 races we can correctly predict and recommend when a user should switch. For Scenario 2, predicting whether a user should keep the assist disabled after they have switched, we have higher precision values compared to Scenario 1 because more information is available to predict a successful disabling.

			Scenario 1: Race before		**Scenario 2: Before & after**	
	N	% success	Prec.	Recall	Prec.	Recall
Stability	23,915	46	.61	.73	.65	.64
Traction	21,066	40	.61	.58	.63	.53
Autobrake	73,244	82	.84	.97	.88	95
ABS	42,535	50	.66	.73	.81	.65
Autoshift	44,314	54	.70	.80	.83	.74
Clutch	11,653	22	(*)	(*)	.63	.29
Full line	62,955	74	.78	.95	.88	.89
Brake line	46,094	51	.68	.76	.83	.68
CosmDmg	31,212	51	.61	.81	.64	.80
LimDmg	20,872	38	.60	.50	.60	.62

Table 5: Precision and recall for Scenarios 1 and 2.
(*) For the clutch assist in Scenario 1 we had a high number of trivial, constant models that we discard from our experiments.

DISCUSSION

Implications for racing game design

Our findings help to improve the built-in assist bundles. For example, patterns of assist usage indicate that the Medium bundle should probably disable the autobrake and the full line, as these assists are the first to be disabled by players, probably because they are the easiest or the least enjoyable. In fact, they are not even found in simpler racing games such as Mario Kart. Moreover, the default Advanced bundle should only enable the traction, stability, and clutch assists.

The damage and brake line assists caught our attention. In online mode, the damage assist is nearly always enabled. Players do collide with each other, but do not want to be penalized for it, so the damage assist should probably be enabled by default in online races. As for the brake line assist, we are puzzled that players maintain it enabled after hundreds of races, especially in online mode. The literature on real-life driving suggests that the faster a driver approaches a turn, the farther ahead he/she needs to look to estimate the curvature of the turn [13], and that experienced drivers look farther ahead than novice drivers [14]. Thus the game could recommend disabling the brake line when switching from a fast to a slow car, and enabling it when the player has not driven the car enough to know the maximum speed at which it can handle a turn. This is especially relevant when racing on an unfamiliar or curvy track, since players may not predict their car's behavior or look far enough when negotiating a difficult turn.

More broadly, although simulation and graphics realism are central in racing games, we believe AI also has a part to play. FM4 provides several levels of AI difficulty in career mode, but does not progressively recommend disabling any assist. Assists should probably remain in the hands of the player, but could the game be more fun if we gave players incentives to disable certain assists? For example, players could unlock an achievement simply for finishing a race with the clutch assist disabled. Moreover, players are relatively predictable in their assist usage, so a recommender system could: 1) accurately nudge players to disable an assist; 2) provide encouragements or reduce the AI difficulty if it believes that the player can make it; and 3) recommend racing with a familiar or low-PI car when disabling an assist for the first time. When we played the game, we found that the career mode adequately ramps up the car difficulty, but it could also be purposed to support the learning curve of certain assists. For example, drag races take place in a straight line with no other car. Integrating them into the career mode could incentivize players to disable the clutch assist.

Game flow

Assists serve multiple purposes. A previous work treats assists as a way to balance the game between players of different skills [6], similar to the notion of handicap in golf. The developers of FM4 somewhat share this vision. Assists do slow down players to make them more consistent and therefore more competitive overall, so they can make the game more fun. But they can also be used as scaffolds when learning. Vygotsky defines the *zone of proximal development* as the tasks a child can do with support but not by herself [15]. A scaffold is a mechanism supporting a learner who has been recently introduced to a new concept. In FM4, the traction and stability assists are scaffolds for learning to drive a RWD car. The line assists are scaffolds for learning a new track. As players progress through races, they learn to race better and disable more and more assists. Taken as such, assists help players learn to race. FM4 highlights that in games, the concepts of fun, zone of proximal development, and flow [16] are interconnected. Game designer Raph Koster argues that a player is in flow when her brain is not entirely able to crunch the game's patterns, and bored when she has mastered all the game's patterns [17]. In FM4, assists are both scaffolds for learning, and tools for fun. But we also show that players may get used to them if they do not disable them early on.

Finally, some of the yoyo players may just have average motor skills and are able to maintain an assist disabled only in certain scenarios. Others may have sufficient motor skills to keep it permanently disabled, but know when to enable the assist to their advantage. On a tortuous track, for example, the former kind of players may enable the brake line to help them anticipate the track. The latter kind of players may enable the brake line to predict the trajectory of the first kind of players, and better pass them. We plan to investigate yoyo players in future work.

CONCLUSION

By analyzing 24 million race entries for 200,000 players, we identified patterns of assist usage in the racing game Forza Motorsport 4. We highlighted noticeable differences between player segments, as well as between assists, and built simple models predicting when a player could disable an assist. Our findings have implications for racing game design and the progression of skill in virtual driving.

ACKNOWLEDGMENTS

Thomas Debeauvais performed this work as an intern at Microsoft Research. We thank Ben Weber and Husain Bengali for supporting this research, and the reviewers for their feedback.

REFERENCES

1 Reeves, S., Brown, B., and Laurier, E. Experts at Play: Understanding skilled expertise. *Games and Culture*, 4, 3 (July 2009), 205-227.

2 Nardi, B. A., Ly, S., and Harris, J. Learning Conversations in World of Warcraft. In *40th Hawaii International Conference on System Sciences* (2007).

3 Huang, J., Zimmermann, T., Nagappan, N., Harrison, C., and Phillips, B. C. Mastering the art of war: how patterns of gameplay influence skill in Halo. In *CHI* (2013), 695-704.

4 Phillips, B. Peering into the Black Box of Player Behavior: The Player Experience Panel at Microsoft Game Studios. In *GDC* (2010).

5 Kraaijenbrink, V., van Gils, F., Cheng, Q., van Herk, R., and van den Hoven, E. Balancing skills to optimize fun in interactive board games. In *Human-Computer Interaction - INTERACT*. 2009.

6 Bateman, S., Mandryk, R. L., Stach, T., and Gutwin, C. Target Assistance for Subtly Balancing Competitive Play. In *CHI* (2011).

7 Hullett, K., Nagappan, N., Schuh, E., and and Hopson, J. Empirical analysis of user data in game software development. In *ESEM* (2012), 89-98.

8 Togelius, J., De Nardi, R., and Lucas, S. M. Towards automatic personalised content creation for racing games. In *Computational Intelligence and Games* (2007), 252-259.

9 Seif El-Nasr, Magy et al., eds. *Game Analytics - Maximizing the Value of Player Data*. Springer, 2013.

10 VGChartz. *Forza Motorsport 4*. http://www.vgchartz.com/game/45656/forza-motorsport-4/.

11 Greenawalt, D. *Online assist*. https://www.youtube.com/watch?v=xMUGP4SKAyg, 2011.

12 Greenawalt, D. *Forza Motorsport 4: Under the Hood Part 3*. http://forzamotorsport.net/en-us/news/underthehood3, 2011.

13 Land, M. and Horwood, J. Which parts of the road guide steering? *Nature*, 377, 6547 (September 1995), 339-340.

14 Milech, D., Glencross, D., and Hartley, L. *Skill acquisition by young drivers: perceiving, interpreting, and responding to the driving environment*. Federal Office of Road Safety, 1989.

15 Vygotsky, L.S. *Mind and society: The development of higher psychological processes*. Harvard University Press, Cambridge, MA, 1978.

16 Csíkszentmihályi, M. *Beyond Boredom and Anxiety: Experiencing Flow in Work and Play*. 1975.

17 Koster, R. *A Theory of Fun for Game Design*. 2004.