Mixed Models as the Basis for Catcher Evaluation and Forecasting

Harry Pavlidis
Baseball Prospectus & Pitch Info

Introduction

Primary Research Team

Harry Pavlidis:

Director of Technology, Baseball Prospectus

Founder, Pitch Info

Dan Brooks:

Contributor, Baseball Prospectus

Lead Scientist, Pitch Info

Jonathan Judge:

Contributor, Baseball Prospectus

Also a lawyer in real life, otherwise he’d be here (I’ll do my best)

Outline

- What is Catcher “Framing”?
- The Value of Framing Revealed
- Regressed Probabilistic Model (“RPM”) - GAM+WOWY
- Called Strikes Above Average (“CSAA”) - MM
- Forecasting Framing

What is Catcher “Framing”?

What is Catcher “Framing”?

The ability to receive a pitch such that the likelihood of a called strike is maximized

Make borderline strikes stay strikes, and steal the occasional ball and turn it into a strike

What is Catcher “Framing”?

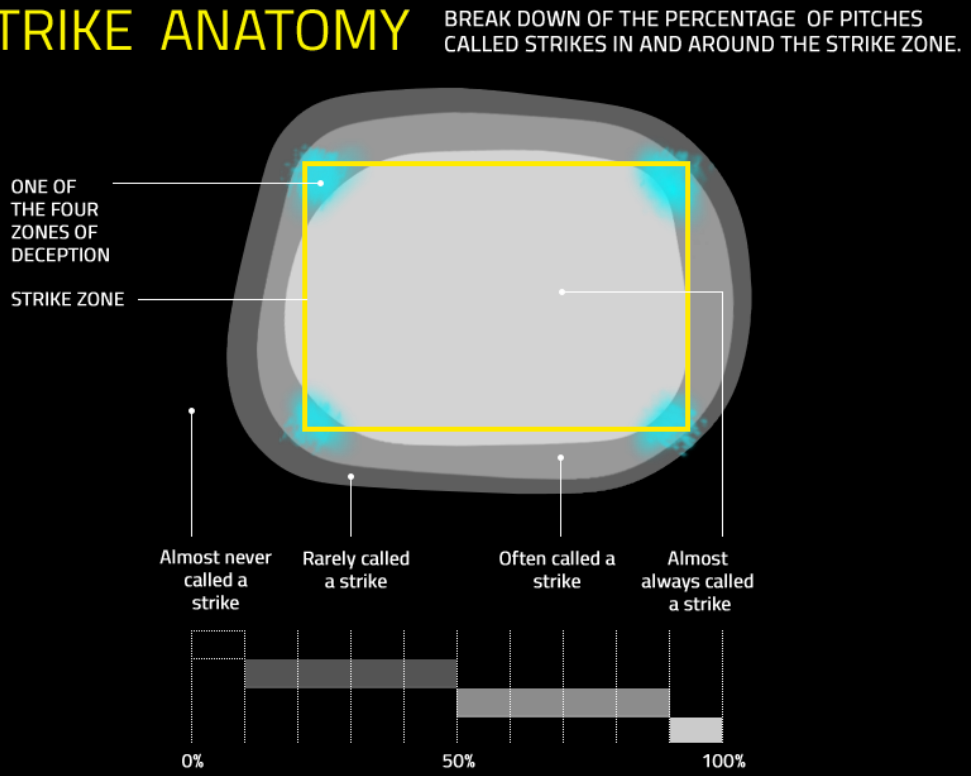

Source: Pavlidis, ESPN The Magazine, July 2014

The Value of Framing Revealed

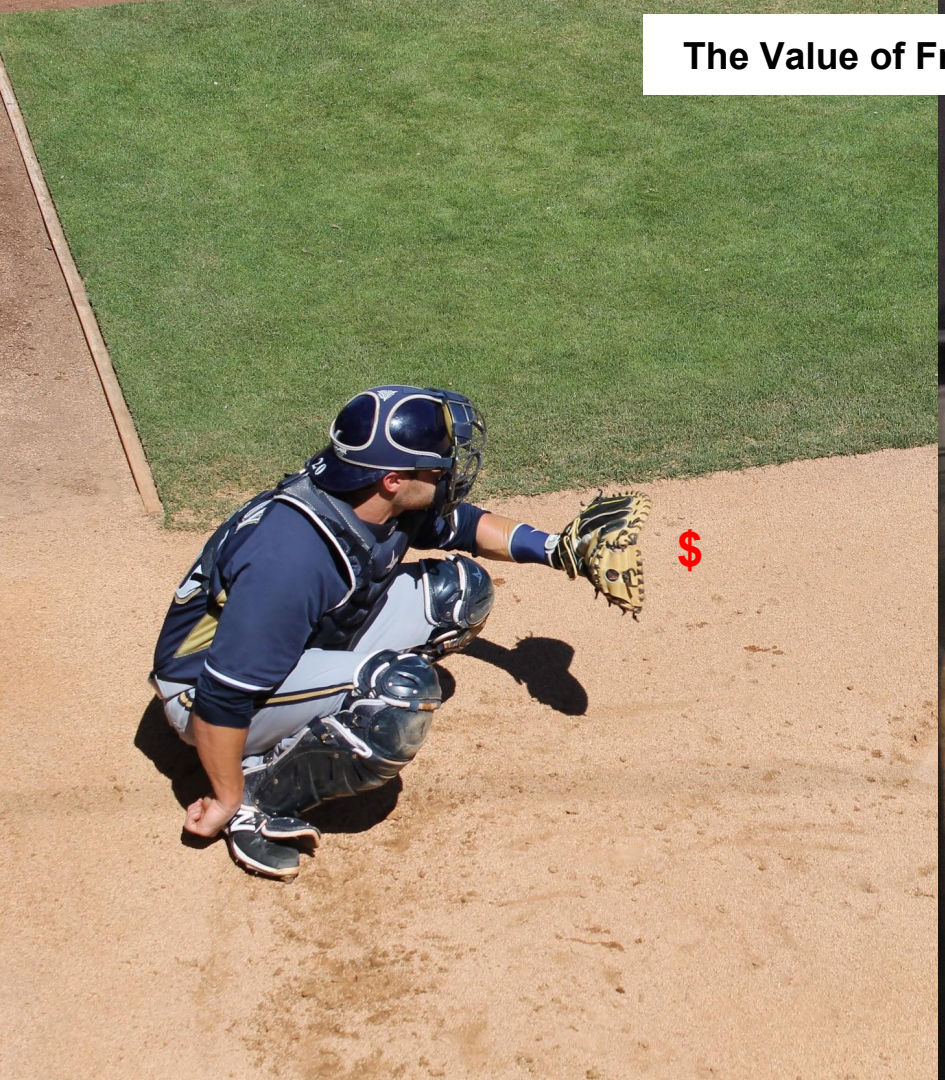

The Value of Framing Revealed

- New Technologies
  - PITCHf/x by Sportvision
    - Data published by MLB Advanced Media
  - Trackman
    - Will replace PITCHf/x
- Research Breakthroughs
  - John Walsh, 2007 (strike zones and values)
  - Dan Turkenkopf, 2008
  - New York Yankees, 2009 (unpublished)
  - Mike Fast, 2011
  - Max Marchi, 2011 & 2013

The Value of Framing Revealed

- Walsh established how the ‘run expectancy’ of an at bat changes from count to count, providing the ‘run value’ of a framed pitch
- Turkenkopf, Fast, Marchi and others found a wide range of skill and value in the MLB catching pool
- Studies were reproduced, and Marchi applied more robust techniques
    - Also tackled historical data that predated tracking technology

The Value of Framing Revealed

Ball	Strike	Run Value
0	0	0.07
0	1	0.08
0	2	0.18
1	0	0.10
1	1	0.10
1	2	0.22
2	0	0.15
2	1	0.15
2	2	0.30
3	0	0.17
3	1	0.24
3	2	0.54

The Value of Framing Revealed

2014

Catcher	Run Value
Buster Posey	27
Rene Rivera	26
Hank Conger	25

Best Seasons 2008-2014

Catcher	Run Value
Jonathan Lucroy 2011	47
Brian McCann 2008	40
Lucroy 2013	35

Total 2008-2014

Catcher	Run Value
McCann	172
Russell Martin	163
Lucroy	159

Source: CSAA model

The Value of Framing Revealed

2014

Catcher	Run Value
Buster Posey	27
Rene Rivera	26
Hank Conger	25

2 - 3 wins

Best Seasons 2008-2014

Catcher	Run Value
Jonathan Lucroy 2011	47
Brian McCann 2008	40
Lucroy 2013	35

3 - 5 wins

Total 2008-2014

Catcher	Run Value
McCann	172
Russell Martin	163
Lucroy	159

18 - 20 wins

Regressed Probabilistic Model using GAM and “WOWY”

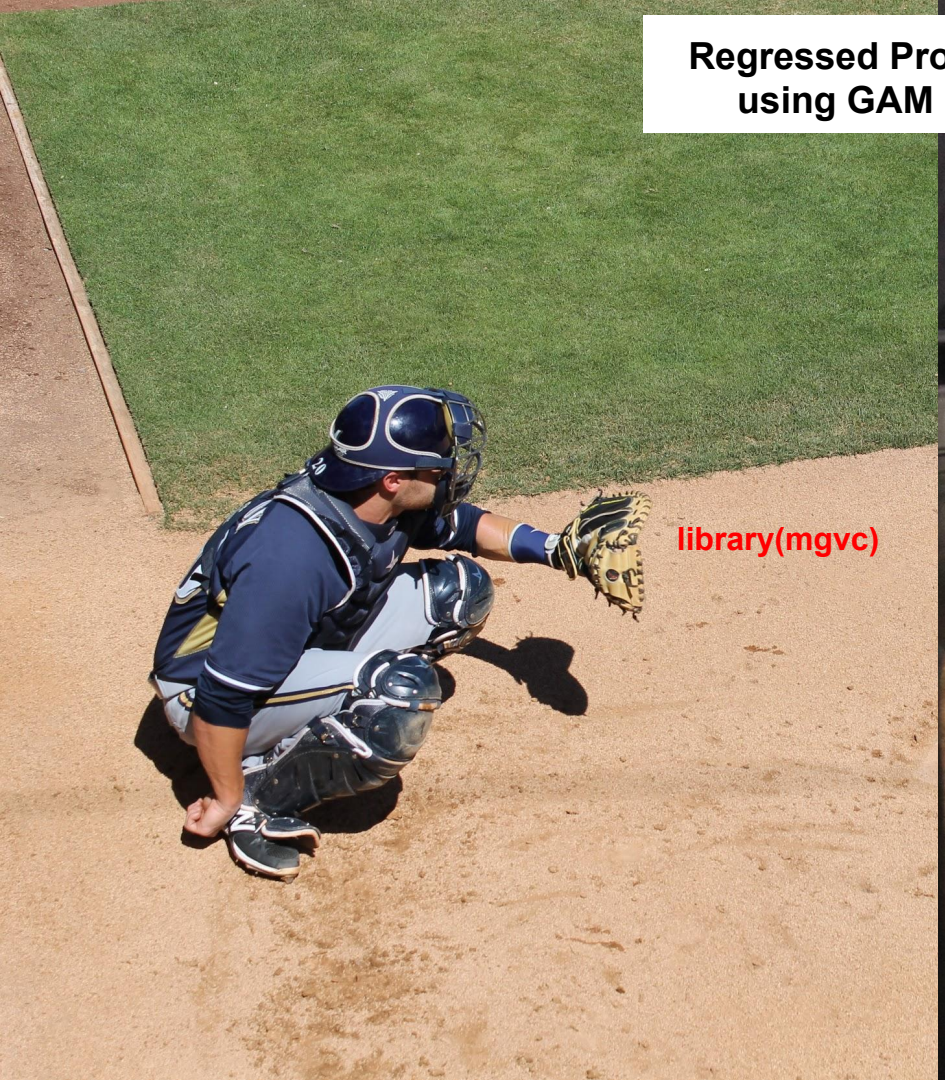

library(mgvc)

RPM: The Impetus

2014

Catcher	Run Value
Buster Posey	27
Rene Rivera	26
Hank Conger	25

2 - 3 wins

Best Seasons 2008-2014

Catcher	Run Value
Jonathan Lucroy 2011	47
Brian McCann 2008	40
Lucroy 2013	35

4 - 6 wins

Total 2008-2014

Catcher	Run Value
McCann	172
Russell Martin	163
Lucroy	159

18 - 20 wins

Didn’t Pass Sniff Tests

RPM: The Goal

Address issues raised by skeptics

- Factoring in pitch type
- Control for batter and pitcher handedness
- Deal with variance in batter strike zones
- Account for the umpire and the pitcher
- Show the numbers!
  - Statcorner publishes +/- numbers
  - The more advanced researchers didn’t share and were hired by teams

RPM: The Data

Pitch Info provides accurate pitch type data and recalibrations to Baseball Prospectus, BrooksBaseball.net and at least one team in every division of Major League Basball.

/appeal to authority

RPM: Generalized Additive Model

For every year, count
and pitch type ...

```
pitchlist$stand<-as.factor(pitchlist$stand)
pitchlist$throws<-as.factor(pitchlist$throws)
pitchlist$ball<-as.factor(pitchlist$ball)
pitchlist$strike<-as.factor(pitchlist$strike)
pitchlist$g_factor <- do.call(paste, c(pitchlist[c("stand","throws")], sep = ""))
pitchlist$g_factor <- as.factor(pitchlist$g_factor)

#model
fit<-bam(c~s(cx,cz,by=g_factor)+ball*strike ,data=pitchlist,family=binomial, cluster=cl)

#predict
myarea <- expand.grid(cx=seq(from=px_min, to=px_max,length.out=lenvar),
cz=seq(from=pz_min, to=pz_max,length.out=lenvar))

#variables from the model
myarea$ball<-bp
myarea$strike<-sp
myarea$g_factor<-paste(standp,throwsp,pg,sep="")

myarea$fit <- c(predict(fit,newdata=myarea,type="response") )

#creates maps for each season/group/throws/stands/count/pitch type
```

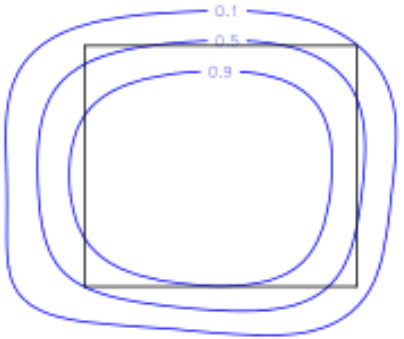

RPM: The Goal Revisited

Address issues raised by skeptics

- Factoring in pitch type
- Control for batter and pitcher handedness
- Deal with variance in batter strike zones
- Account for the umpire and the pitcher
- Show the numbers!
  - Statcorner publishes +/- numbers
  - The more advanced researchers didn’t share and were hired by teams

RPM: With or Without You ("WOWY")

Developed by Tom Tango http://www.tangotiger.net/catchers.html

WOWY compares all the available pitcher/catcher combinations in order to assess whether a particular catcher, given the pitchers he has caught throughout his career, has fared better or worse than expected. - Marchi, Hardball Times, 2012

RPM created adjustments based on WOWY for each pitcher and umpire and applied that to each pitch called.

catcher_score = call - CS_probability - pitcher_adjustment - umpire_adjustment

call = 1 for strike, 0 for ball; multiply score by count-specific or generic run value for total value

RPM: Regression

- Regress each catcher to league average (0)
- Regression was based on intraclass correlation ("ICC") of each measurement.
- ICC consistency and agreement both showed that a 50/50 point occurred after ~290 framing opportunities

The amount to regress “r” is based on the framing chances “n” and this 50/50 point

r = n / (n + 290)

RPM: Agreement With Scouting

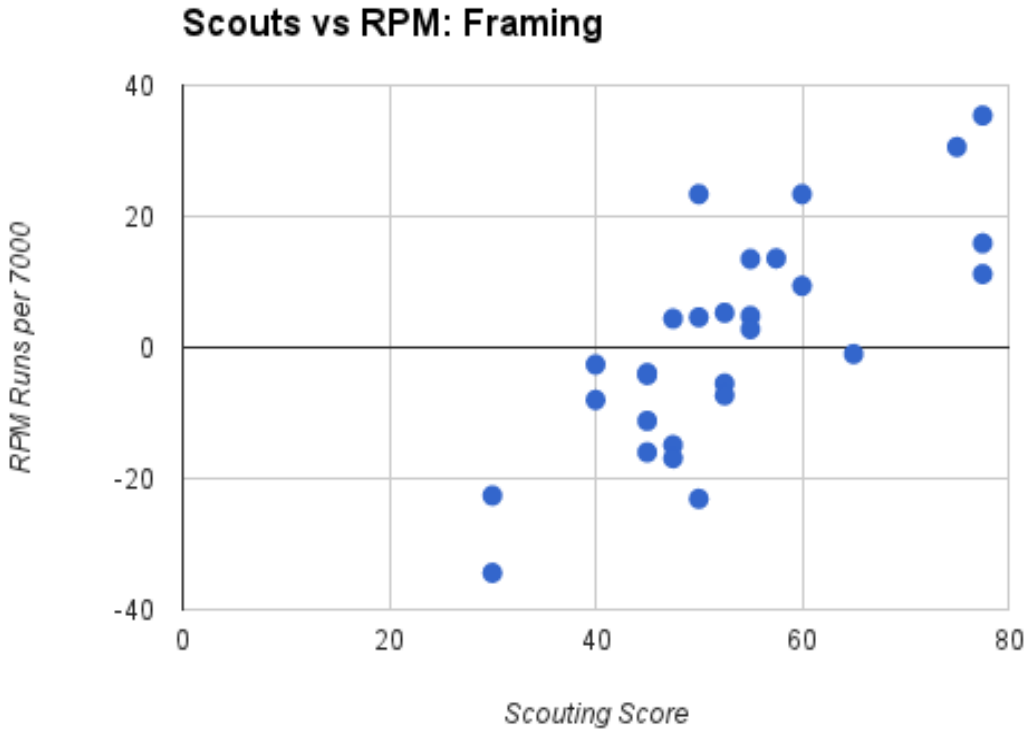

RPM: Example 1, Red Sox Upgrade

Christian Vazquez

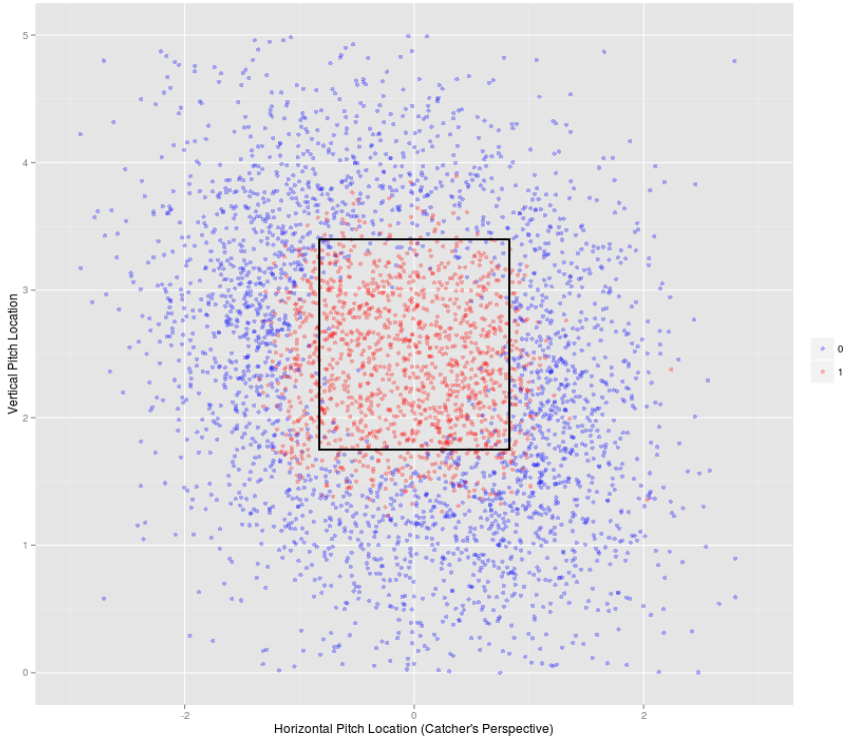

A.J. Pierzynski

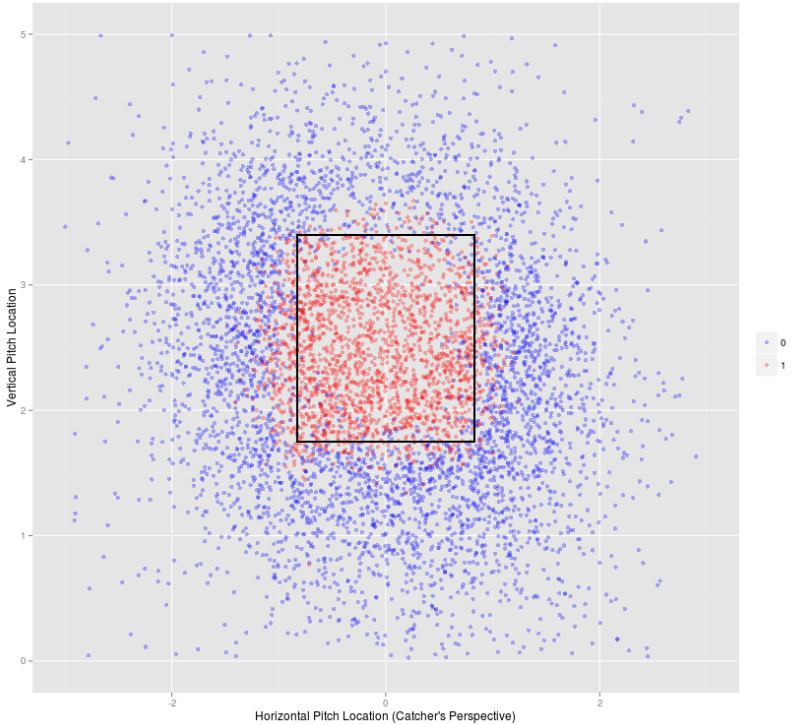

RPM: Example 2, The Extremes

Wilin Rosario

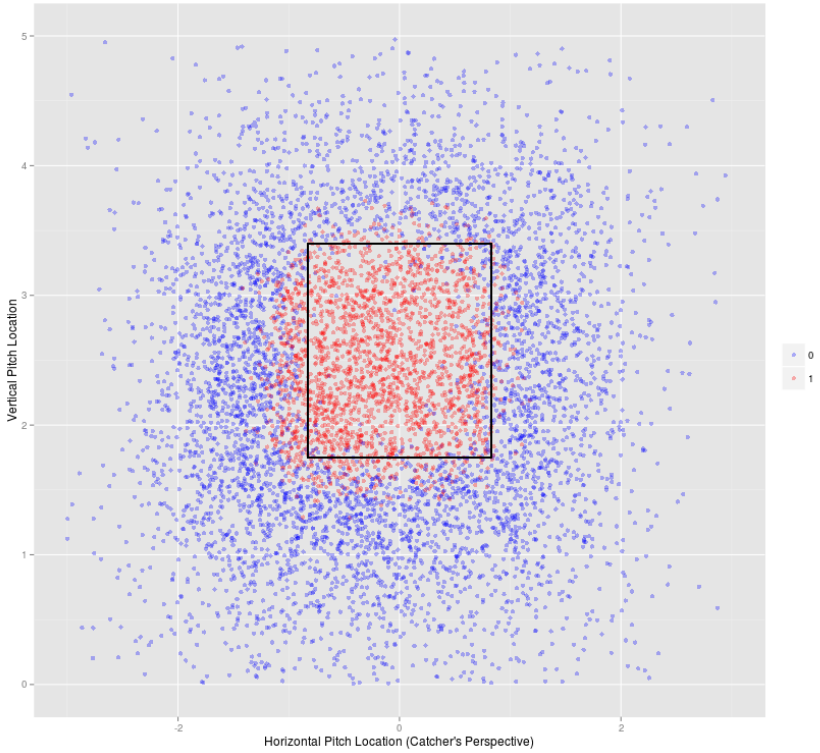

Jonathan Lucroy

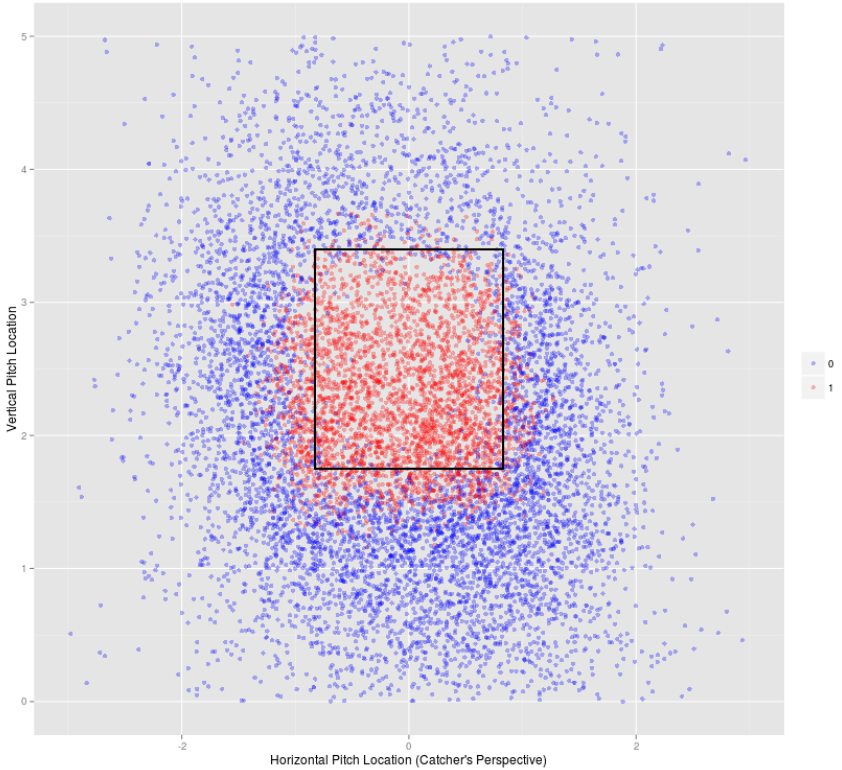

RPM: We Published It

BIOGRAPHICAL

edit bio info

Jose Molina C Rays

Rays Player Cards | Rays Team Audit | Rays Depth Chart

2014 PROJECTIONS (SEASONAL AGE 39)

PA	AVG	HR	R	RBI	SB	TAv	WARP
414	.223	6	37	35	3	.227	-0.5

Birth Date	6-3-1975	Age	38 years, 8 months, 27 days
Height	6' 2"	Bats	R
Weight	250 lbs	Throws	R

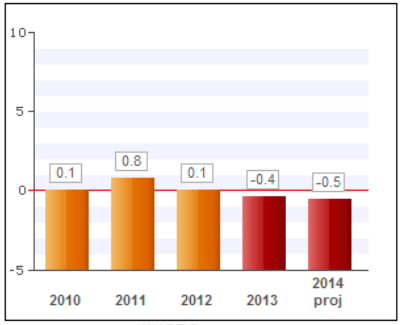

WARP Summary

SHOW ALL | STATS ONLY | TRANSACTIONS | SPLITS | BOOK COMMENTS | PITCH F/X | ARTICLE | CATCHING

Extra Catcher Data by Year

Year	Framing Chances	Extra Strikes	Count-Sensitive Framing Runs	Framing Runs per 7000	Blocking Chances	PB/WP Saved	Blocking Runs	Blocking Runs per 7000	Total Receiving Runs	Total Receiving Runs per 7000
2008	5384	206	25.6	33.3	6692	-23.4	-6.6	-6.9	19.1	26.5
2009	2708	156	19.1	49.3	3462	-10.6	-3.0	-6.0	16.1	43.3
2010	3203	151	18.7	40.8	4436	-22.5	-6.3	-9.9	12.4	30.8
2011	1825	59	6.9	26.6	3824	-16.9	-4.7	-8.6	2.2	17.9
2012	4747	192	24.0	35.4	6378	-27.0	-7.5	-8.3	16.5	27.1
2013	4786	174	23.9	34.9	6732	-24.8	-6.9	-7.2	17.0	27.7
total	22653	930	116.3	35.9	31524	-117.2	-32.8	-7.3	83.5	28.6

A Collaboration between BrooksBaseball.net and Baseball Prospectus - Pitch classifications provided by Pitch Info LLC

Called Strikes Above Average: Mixed Models

library(lme4)

CSAA: The Goal

- Replace WOWY and ‘manual’ regression with a more robust technique
- Replicate Marchi’s “retrosheet” work
- Make Judge do most of the work
- Publish it all

CSAA: Why a Mixed Model?

Fixed Effects

- called strike probability (from RPM's GAM)
- ball/strike
- home field advantage

Random Effects

- catcher
- pitcher
- umpire
- batter

CSAA: PITCHf/x

IV selection driven by theory and BIC

The null is the hardest part …

Final value …

```
y.model  <-    glmer(cs ~ half+ cs_prob + (1| catcher) + (1| batter)
+  (1|pitcher) + (1|umpire) + (1|catcher:cs_prob), y,
family=binomial(link = "probit"), nAGQ = 0)#

y.model.f <-  fixef(y.model)
y.model.r <-  ranef(y.model)

catcher.csprob.interact <- y.model.r$"catcher:cs_prob"
catcher.r <- y.model.r$catcher
pitcher.r <- y.model.r$pitcher
umpire.r <- y.model.r$umpire
batter.r <- y.model.r$batter

##calculate null effects
fixed.int <- as.numeric((y.model.f[1])  )


null <- fixed.int
+ median(y$cs_prob)
+ median(catcher.r$int)
+ median(catcher.csprob.interact$int)
+ median(umpire.r$int)
+ median(pitcher.r$int)
+ median(batter.r$int)

null.prob <- pnorm(null)

 catcher.r$sum.int <- (catcher.r$int +  null)
 catcher.r$prob <- pnorm(catcher.r$sum.int)
 catcher.r$CSAA <- catcher.r$prob - null.prob
```

CSAA: Historical

IV selection driven by theory and BIC

The null is the hardest part …

Final value …

```
y.model <- glmer(cs ~ stands*throws +  home_batting + count +
(1|catcher) + (1|batter) + (1|umpire) + (1|pitcher), y,
family=binomial(link = "probit"), nAGQ = 0)

y.model.f <-  fixef(y.model)

y.model.r <-  ranef(y.model)

catcher.r <- y.model.r$catcher
pitcher.r <- y.model.r$pitcher
umpire.r <- y.model.r$umpire
batter.r <- y.model.r$batter

##calculate null effects
fixed.int <- as.numeric(y.model.f$int  )


null <- fixed.int
+ median(catcher.r$int)
+ median(umpire.r$int)
+ median(pitcher.r$int)
+ median(batter.r$int)

null.prob <- pnorm(null)

catcher.r$PLAYERID <- rownames(catcher.r)
rownames(catcher.r) <- NULL
catcher.r$sum.int <- (catcher.r$int +  null)
catcher.r$prob <- pnorm(catcher.r$sum.int)
catcher.r$CS.Prob <- catcher.r$prob - null.prob
```

Forecasting Framing

Forecasting: PECOTA or Not

Projection system developed by Nate Silver and Clay Davenport and maintained by Baseball Prospectus employee Rob McQuown.

Relies on pools of comparable players, minor league equivalents (MLE) and aging curves to generate a range of projections for each player.

Projects Fielding Runs Above Average (“FRAA”)

We can add CSAA to FRAA

For now let’s project CSAA alone...

Forecasting Framing: Challenges

- Limited player pool
- Rule changes
  - strike zone enforcement
  - actual change to strike zone definition
- Umpire training
  - they use PITCHf/x
- Limited Minor League data

Forecasting Framing: Aging

We may lack enough variance for fielding comps, but we have a good idea of how catchers age from 25-35

Catching seasons dating to 1988 were used

Results consistent with Marchi's study

Survivor bias: generate a conservative projection for catchers who don't continue (approx -0.003 CSAA)

```
aging <-
      ddply(d.age.surv, c("age.x"), summarise,
      gap = as.numeric(wtd.mean(gap,weights= n.harmonic)) ,
       n = length(age.x)       )
aging$cumu.gap <- cumsum(aging$gap)
fit <- loess(  aging$cumu.gap~aging$age.x,weights=aging$n ,span=.55)
```

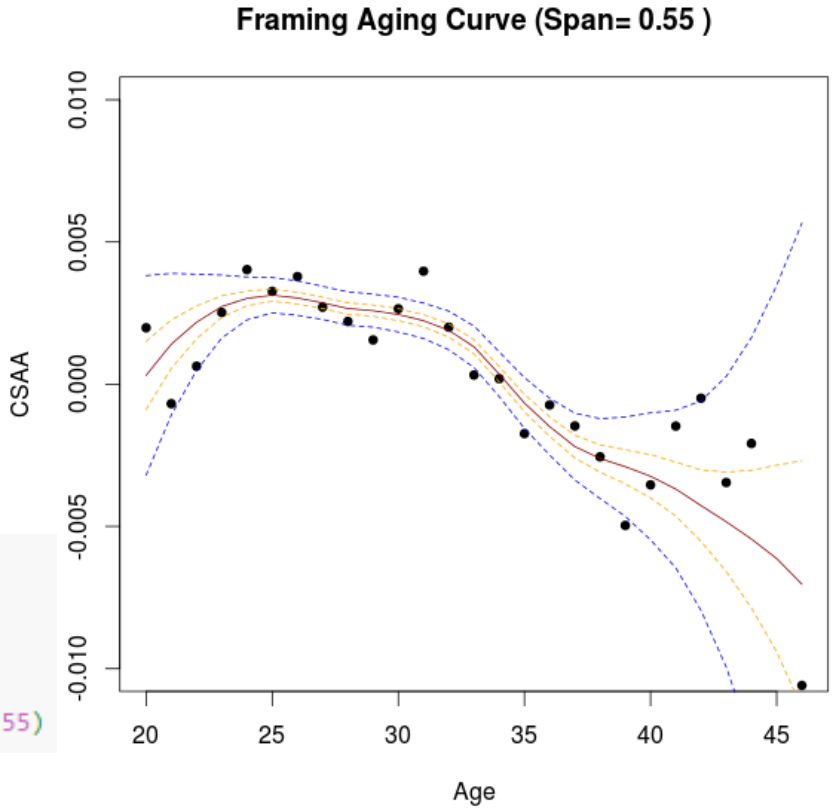

Forecasting Framing: Testing 2014

- Use the basic 5|3|1 * chances for weighting for past seasons
- Age the weighted value
- Hope for the best

Weighted Correlation r=.797 stderr=.064

RMSE .008

MAE .007

*Std Dev of CSAA ~ .10 to .12

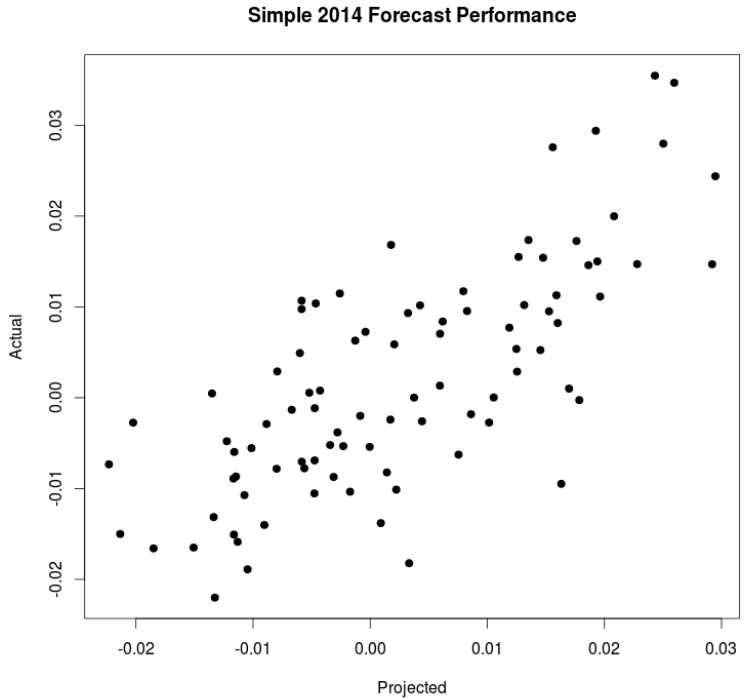

Forecasting Framing: 2015

Setting each to a 7000 chance season

Catcher	Run Value
Rene Rivera (recently traded to Tampa)	32
Hank Conger (recently traded to Houston)	30
Christian Vazquez (promoted in 2014 by Boston)	28

Thank You