Isolating Individual Skater Impact on Team Shot Quantity and Quality

Micah Blake McCurdy
hockeyviz.com

Ottawa, Canada
Ottawa Hockey Analytics Conference
September 15, 2018

Introduction

Tired: Numbers

Wired: Pictures

Introduction

Tired: Numbers

Wired: Pictures

With Evgeni Malkin and Connor McDavid on the ice at 5v5 during 2017-2018, their teams generated 49 and 51 unblocked shots per hour, respectively; that is, 13% and 18% more than league average.

Malkin & McDavid *On-Ice* Offence

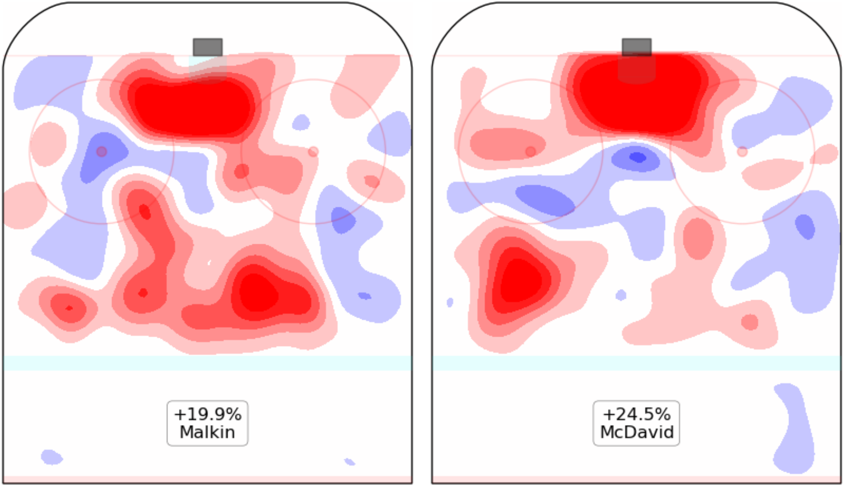

Malkin & McDavid *On-Ice* Defence

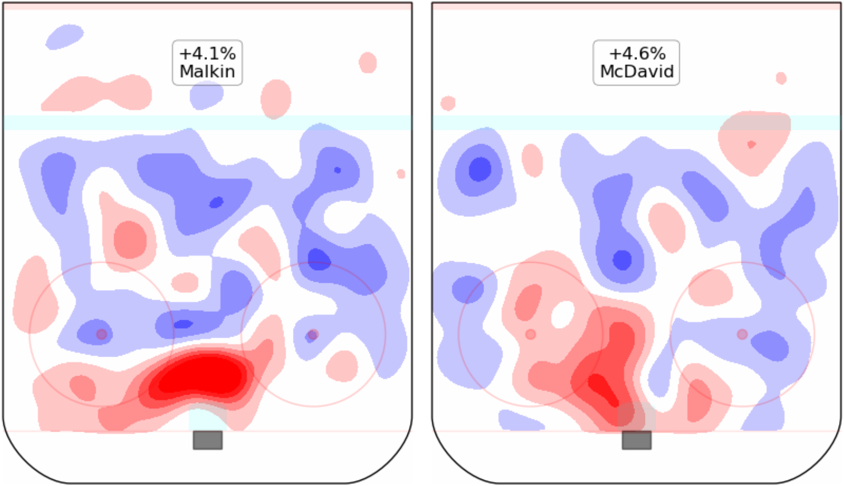

Aim

Isolate individual skater impact on team shots, both for and against.

New Thing

Treat maps as first-class objects, instead of single-numbers like rates or counts.

Isolation

Control for the most important aspects of play which are *outside* of a player's control:

- Other skaters
  - Teammates
  - Opponents
- Zone usage
- The score (slyly sneaking in coaching, maybe)

Isolation

Control for the most important aspects of play which are *outside* of a player's control:

- Other skaters
  - Teammates
  - Opponents
    - Gonna settle this once and for all. For all!!
- Zone usage
- The score (slyly sneaking in coaching, maybe)

Bayesian Approach

Begin with an extremely simple estimate of player ability and update it slowly after every shot.

Bayesian Approach

Begin with an extremely simple estimate of player ability and update it slowly after every shot.

Before the observations begin, what do we know about the players?

Bayesian Approach

Begin with an extremely simple estimate of player ability and update it slowly after every shot.

Before the observations begin, what do we know about the players?

They're all NHL players.

Bayesian Approach

Begin with an extremely simple estimate of player ability and update it slowly after every shot.

Before the observations begin, what do we know about the players?

They're all NHL players.

Take as prior that every player is a league average player.

Bayesian Approach

Begin with an extremely simple estimate of player ability and update it slowly after every shot.

Before the observations begin, what do we know about the players?

They're all NHL players.

Take as prior that every player is a league average player. (With some sneakiness about players with very little icetime.)

Dessert First

How I did it:

Dessert First

How I did it: Later (math was implicated)

Dessert First

How I did it: Later (math was implicated)
For now, what is the good of it?

Dessert First

How I did it: Later (math was implicated)
For now, what is the good of it?

Dessert First

How I did it: Later (math was implicated)
For now, what is the good of it?

- The base layer of a "full-value" model (a war-like stat).

Dessert First

How I did it: Later (math was implicated)

For now, what is the good of it?

- The base layer of a “full-value” model (a war-like stat).
    - Not interesting to me for the foreseeable future

Dessert First

How I did it: Later (math was implicated)
For now, what is the good of it?

- The base layer of a "full-value" model (a war-like stat).
  - Not interesting to me for the foreseeable future
- The base layer of a set of models to simulate hockey games.

Dessert First

How I did it: Later (math was implicated)
For now, what is the good of it?

- The base layer of a “full-value” model (a war-like stat).
  - Not interesting to me for the foreseeable future
- The base layer of a set of models to simulate hockey games.
  - Yes ok sure

Dessert First

How I did it: Later (math was implicated)
For now, what is the good of it?

- The base layer of a “full-value” model (a war-like stat).
  - Not interesting to me for the foreseeable future
- The base layer of a set of models to simulate hockey games.
  - Yes ok sure
- Understand which players are victims of circumstance and which the beneficiaries.

Dessert First

How I did it: Later (math was implicated)
For now, what is the good of it?

- The base layer of a "full-value" model (a war-like stat).
  - Not interesting to me for the foreseeable future
- The base layer of a set of models to simulate hockey games.
  - Yes ok sure
- Understand which players are victims of circumstance and which the beneficiaries.
  - Extremely satisfying but still one-dimensional.

Dessert First

How I did it: Later (math was implicated)
For now, what is the good of it?

- The base layer of a “full-value” model (a war-like stat).
  - Not interesting to me for the foreseeable future
- The base layer of a set of models to simulate hockey games.
  - Yes ok sure
- Understand which players are victims of circumstance and which the beneficiaries.
  - Extremely satisfying but still one-dimensional.
- Use it to see how different players affect how offence/defence moves through different parts of the ice.

Dessert First

How I did it: Later (math was implicated)
For now, what is the good of it?

- The base layer of a “full-value” model (a war-like stat).
  - Not interesting to me for the foreseeable future
- The base layer of a set of models to simulate hockey games.
  - Yes ok sure
- Understand which players are victims of circumstance and which the beneficiaries.
  - Extremely satisfying but still one-dimensional.
- Use it to see how different players affect how offence/defence moves through different parts of the ice.
  - VERY YES

Dessert First

How I did it: Later (math was implicated)
For now, what is the good of it?

- The base layer of a “full-value” model (a war-like stat).
  - Not interesting to me for the foreseeable future
- The base layer of a set of models to simulate hockey games.
  - Yes ok sure
- Understand which players are victims of circumstance and which the beneficiaries.
  - Extremely satisfying but still one-dimensional.
- Use it to see how different players affect how offence/defence moves through different parts of the ice.
  - VERY YES
  - With tracking data we could work even more baroque and byzantine things into the same framework

Threat

To form summary statistics we weight shot maps according to league average shooting percentages from given locations to obtain *threat*.

Threat

To form summary statistics we weight shot maps according to league average shooting percentages from given locations to obtain *threat*.

Carefully avoiding shooting talent and goaltender talent.

Threat

To form summary statistics we weight shot maps according to league average shooting percentages from given locations to obtain *threat*.

Carefully avoiding shooting talent and goaltender talent.

- Units of threat are goals per hour;
- Threat is like the worst xG model that is still worth writing down.

Malkin and McDavid Threat

With Evgeni Malkin and Connor McDavid on the ice at 5v5 during 2017-2018, their teams threatened 2.8 and 2.9 goals per hour, respectively; that is, 20% and 25% more than league average.

On-ice Observed Threat

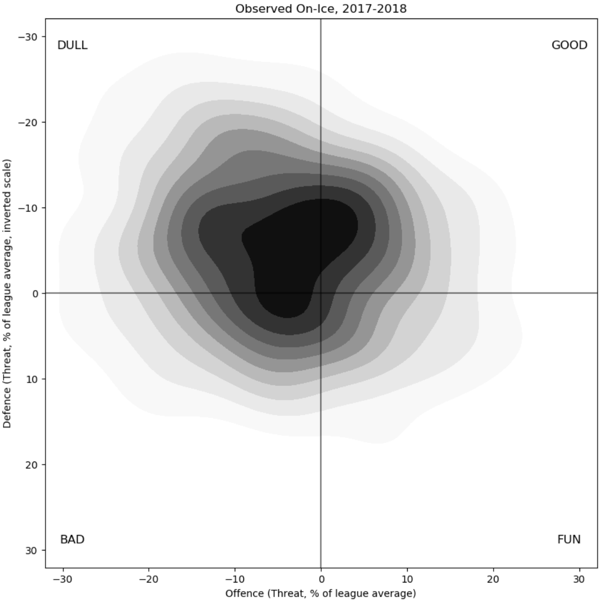

On-ice Observed Threat

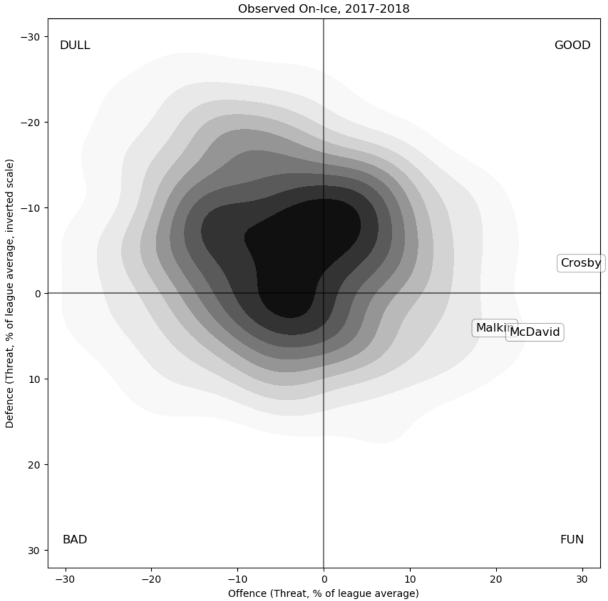

Teammates Distribution

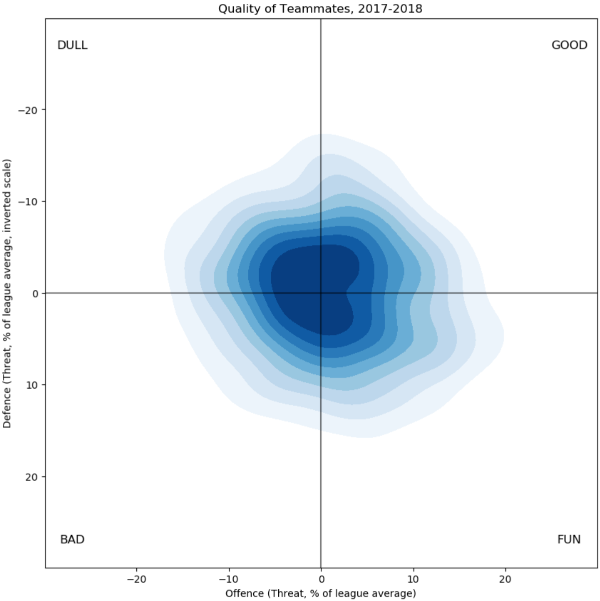

Teammates Distribution

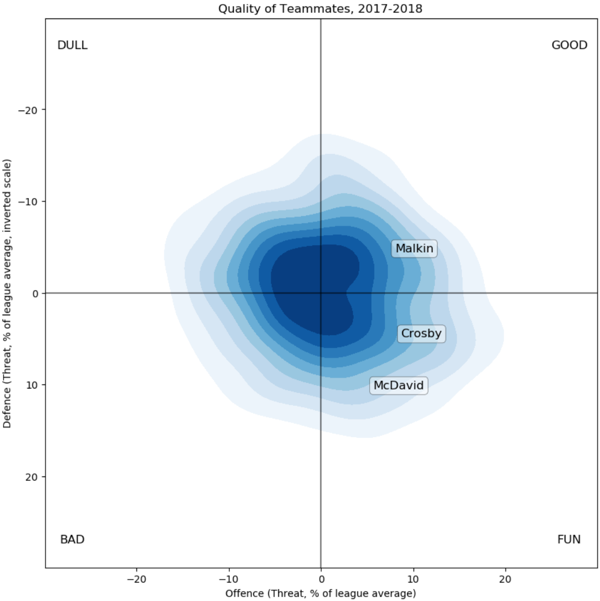

Competition Distribution

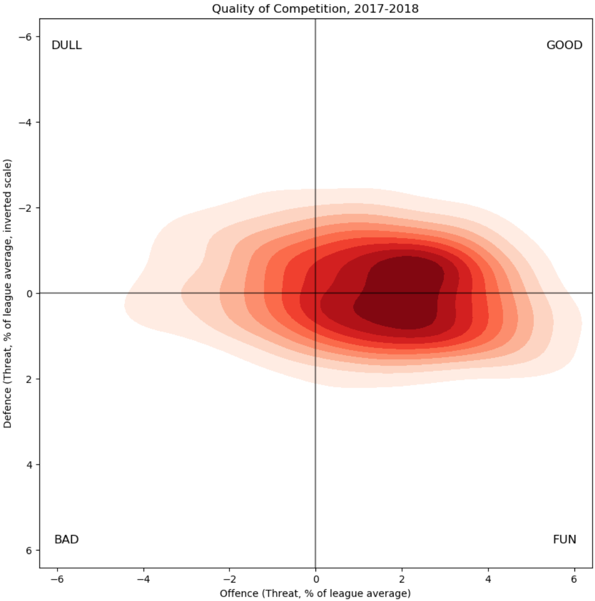

Competition Distribution

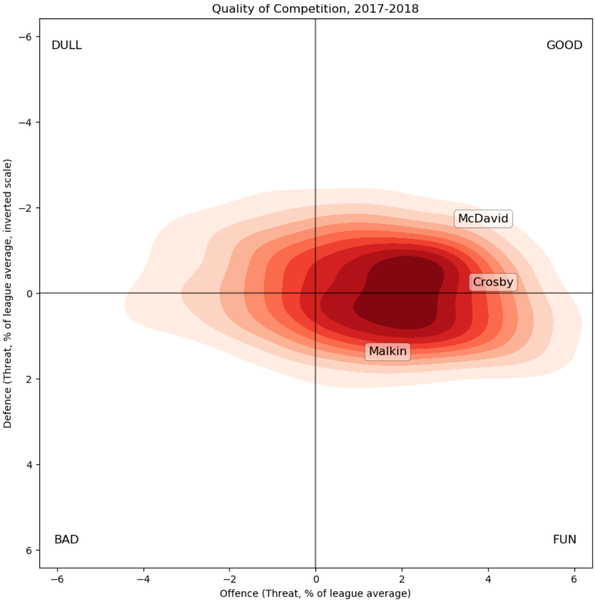

Score Impact

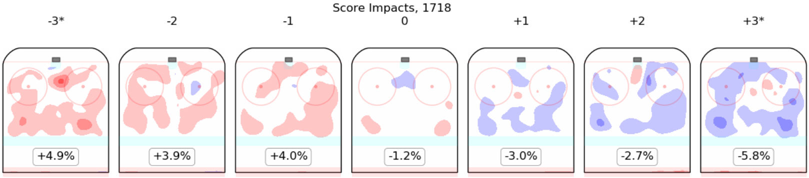

Score Distribution

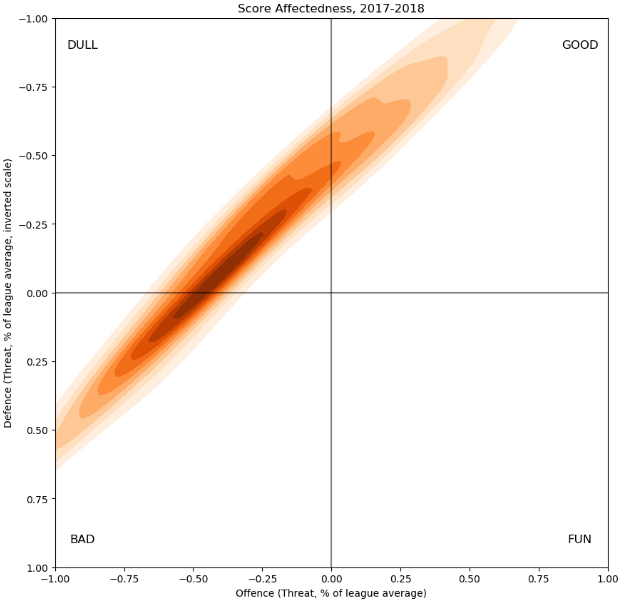

Score Distribution

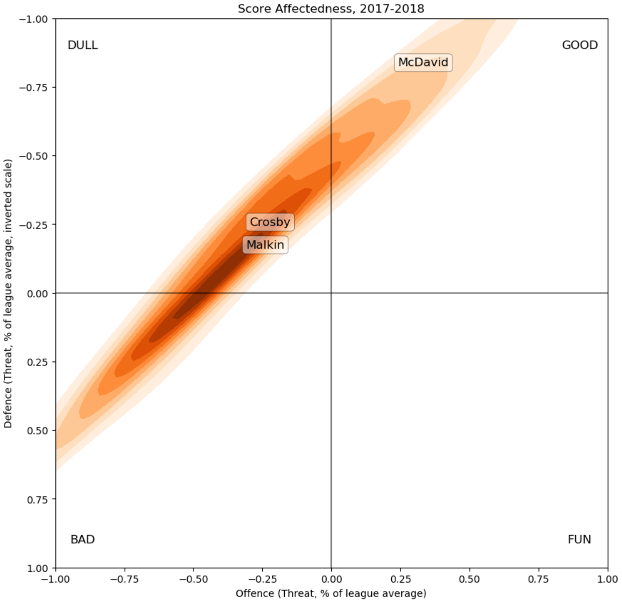

Zone Impact (Home Teams)

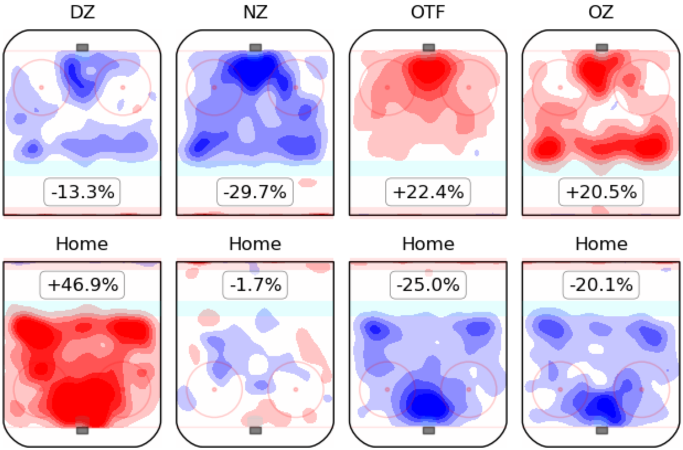

Zone Impact (Away Teams)

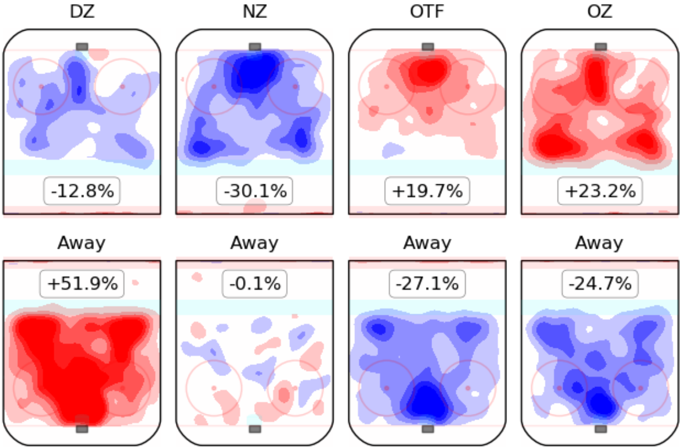

Zone Distribution

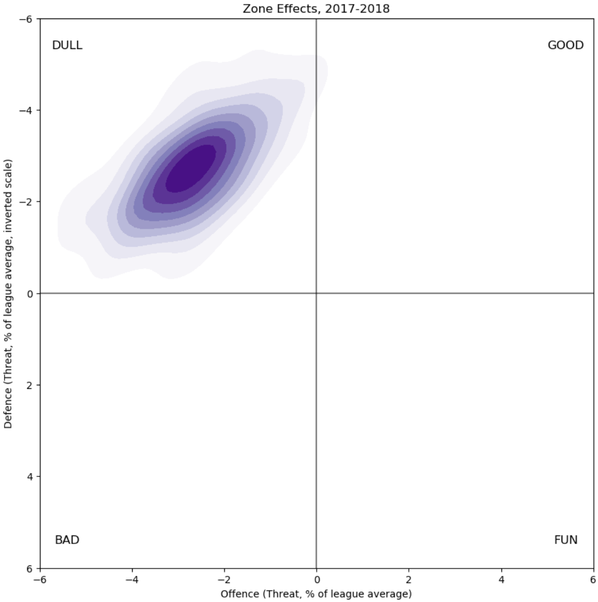

Zone Distribution

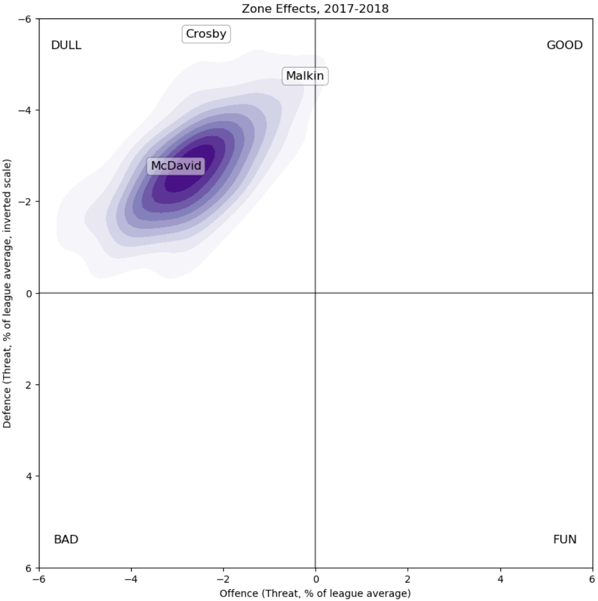

"Residue" Distribution

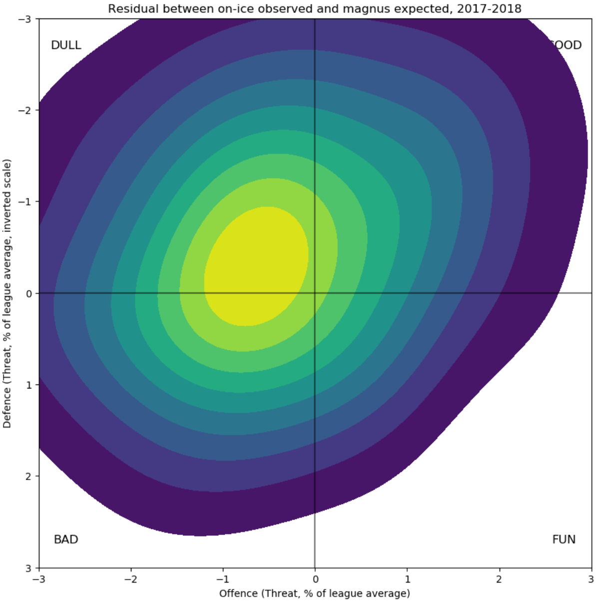

"Residue" Distribution

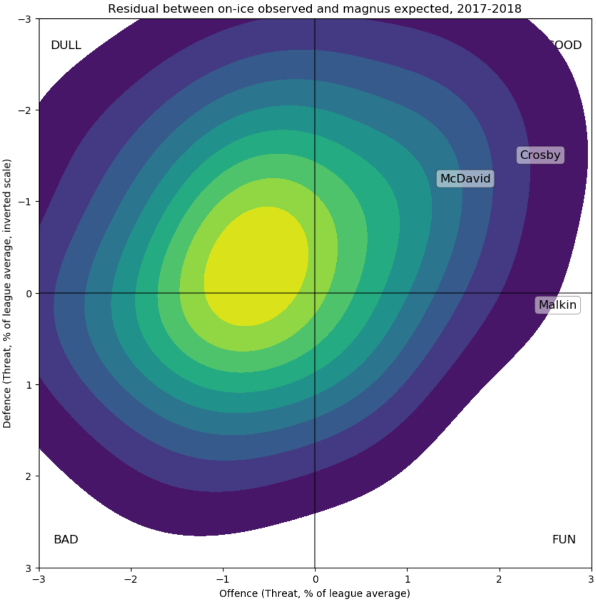

Isolated Individual Impact

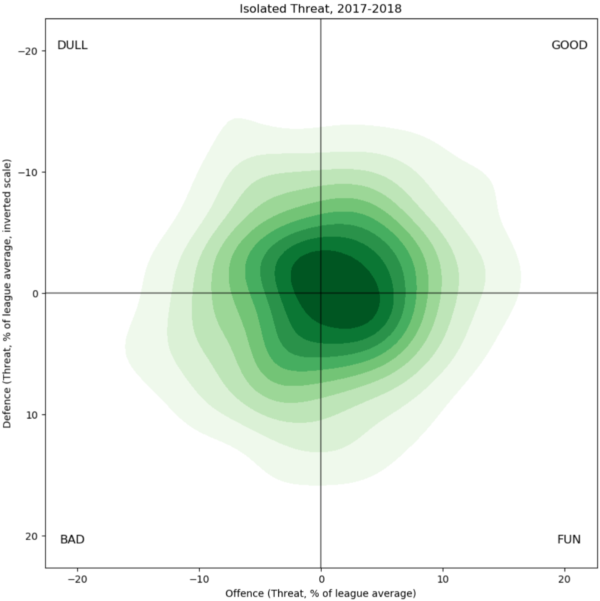

Isolated Individual Impact

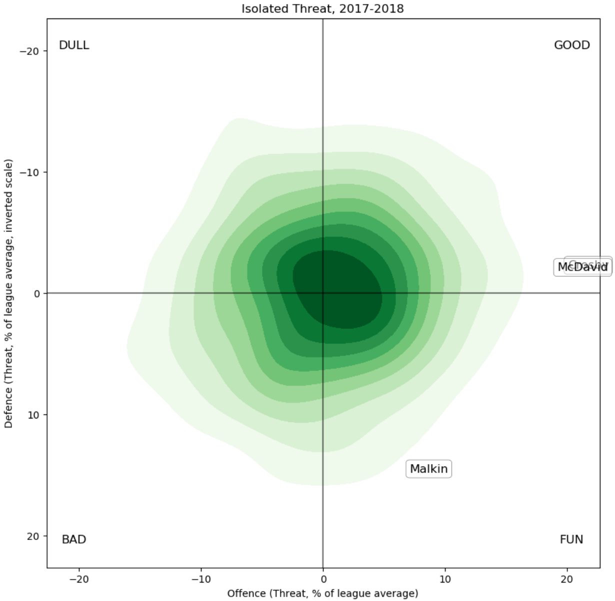

Malkin & McDavid *Isolated* Offence

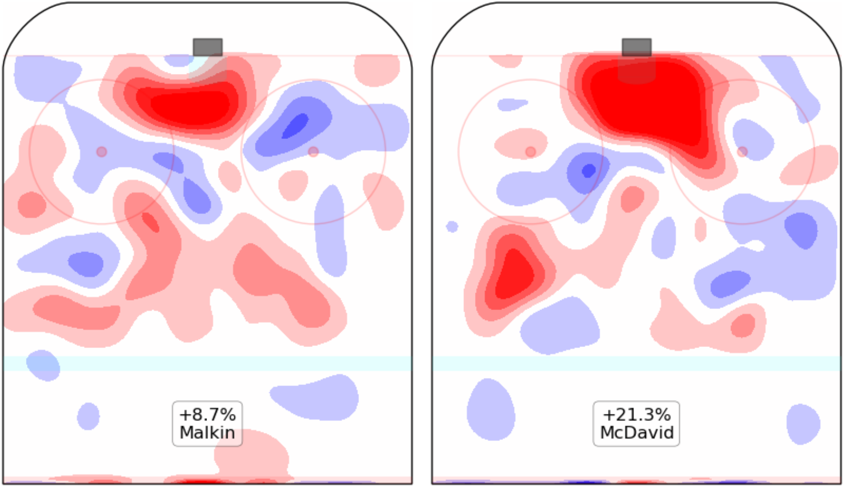

Malkin & McDavid *Isolated* Defence

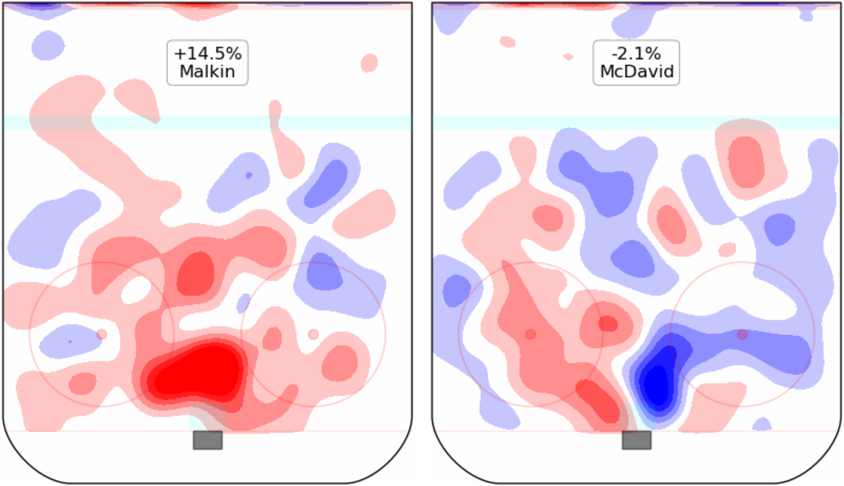

Scale

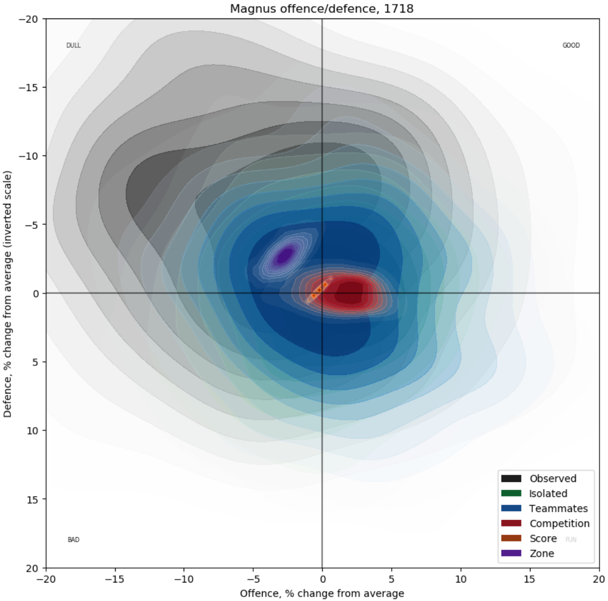
Magnus offence/defence, 1718
DULL
GOOD
BAD
FUN
Defence, % change from average (inverted scale)
Offence, % change from average
−20
−15
−10
−5
0
5
10
15
20
Observed
Isolated
Teammates
Competition
Score
Zone

Scale

As a whole, teammates have a much larger impact than competition; about five times as much. For some individual players; competition impact is still larger than teammate impact.

Competition Dominates Context for Some Players

For instance:

Player	Team	Teammate Impact	Competition Impact
Brandon Saad	CBJ	+0.2%	-6.4%
Nazem Kadri	TOR	-0.2%	-4.7%
Chris Thorburn	STL	+0.5%	+4.6%

Method

Bayesian update as vaguely gestured at earlier can be implemented by (generalized) *"ridge" regression*. (Totally unrelated to the motivations of the people who first suggested it for totally technical reasons in situations not at all resembling ours, because math.)

Method

Bayesian update as vaguely gestured at earlier can be implemented by (generalized) *"ridge" regression*. (Totally unrelated to the motivations of the people who first suggested it for totally technical reasons in situations not at all resembling ours, because math.) Take a linear model of the form

$$Y = X\beta$$

where:

- Y is what you see on the ice.
- X is the design matrix (players (twice), zones, scores, intercepts)
- β is the estimates of the impact of each model feature.

Design

- *Columns* of X correspond to model features - things that we imagine affect what happens. There are around $2,000$.
- *Rows* of X correspond to slivers of hockey where none of those things change - "microshifts". For a single season, around a million.
- *Values* in X are almost entirely indicators (zeros or ones).

Units of β and Y

The only thing we need for β (our estimate of ability/impact) and Y (our "observations" of what happened on the ice) is to do is:

- *have the same units*;
- *be things that can be added together and multiplied by numbers* and;
- *be things that have some notion of "size"*.

(In fact any inner product space will do)

Units of β and Y

The only thing we need for β (our estimate of ability/impact) and Y (our "observations" of what happened on the ice) is to do is:

- *have the same units*;
- *be things that can be added together and multiplied by numbers* and;
- *be things that have some notion of "size"*.

(In fact any inner product space will do)

- Like a shot rate density map

Units of β and Y

The only thing we need for β (our estimate of ability/impact) and Y (our "observations" of what happened on the ice) is to do is:

- *have the same units*;
- *be things that can be added together and multiplied by numbers* and;
- *be things that have some notion of "size"*.

(In fact any inner product space will do)

- Like a shot rate density map
  - Or a distribution of shot rate density maps...
  - Or a distribution of trajectories through a configuration space of all possible ways twelve hockey players can be placed and oriented on a hockey rink ....

How to *fit* the model

Given a model of the form $Y = X\beta$, where we know X and Y, how do we find β?

How to *fit* the model

Given a model of the form $Y = X\beta$, where we know X and Y, how do we find β?
We *could* simply find the β with the smallest (total) deviation from the observations, like a chump. That is, minimize

$$(Y - X\beta)^T(Y - X\beta)$$

which is solved by

$$\beta = (X^TX)^{-1}X^TY$$

For hockey, this will usually be overfit; that is, it will follow the chaos and noise of the data very closely (much too closely).

How to *fit* the model

Given a model of the form $Y = X\beta$, where we know X and Y, how do we find β?
Instead, we could use our assumption that the players are NHL players and instead minimize:

$$\text{Error} = (Y - X\beta)^T (Y - X\beta) + \beta^T \Lambda \beta$$

where Λ is a matrix which encodes our prior information that the players are all NHL players. This is *zero-biased* regression where we set our zero at league average.
Solved by:

$$\beta = (X^T X + \Lambda)^{-1} X^T Y$$

Tuning

We get to choose Λ since it encodes our *prior* information, before we examine the observations from the season. Choosing suitable values for its entries is a matter of skill and artifice (there is some guessing and eyeballing).

Much scope for very subtle priors if we want; I use a diagonal Λ with entries:

- $\lambda = 10,000$ for all players and zones and scores, and
- $\lambda = 0.001$ for the intercepts, except:
- λ varying from $2,000$ to $10,000$ for low ice-time players.

Crowd Pleasers (Guys recently in strange situations)

Who was best/worst/weirdest in the last season?

Best 5v5 Offensive Threat Performances, 2017-2018

(thousand minute minimum)

Player	Team	Isolated Threat Relative to league average
Sidney Crosby	PIT	+21.7%
Connor McDavid	EDM	+21.3%
Roman Josi	NSH	+12.1%
John Klingberg	DAL	+11.5%
Kris Letang	PIT	+9.7%
Drew Doughty	L.A	+8.7%
Jeff Petry	MTL	+8.3%
Alex Radulov	DAL	+7.9%
Artemi Panarin	CBJ	+7.7%
Marc-Edouard Vlasic	S.J	+7.6%

Best 5v5 Defensive Threat Performances, 2017-2018

(thousand minute minimum)

Player	Team	Isolated Threat Relative to league average
Mikko Koivu	MIN	-17.9%
Greg Pateryn	DAL	-16.0%
Evgeni Dadonov	FLA	-14.4%
Hampus Lindholm	ANA	-12.9%
Carl Hagelin	PIT	-12.1%
Colton Parayko	STL	-12.0%
Radko Gudas	PHI	-10.0%
Brayden Point	T.B	-10.0%
Alex Iafallo	L.A	-9.9%
Niklas Kronwall	DET	-9.6%

Best 5v5 Net Threat Performances, 2017-2018

(thousand minute minimum)

Player	Team	Isolated Threat Relative to league average
Evgeni Dadonov	FLA	+25.9%
Sidney Crosby	PIT	+24.0%
Connor McDavid	EDM	+23.4%
Brendan Gallagher	MTL	+23.4%
Pierre-Luc Dubois	CBJ	+23.1%
Colton Parayko	STL	+17.6%
Jordan Eberle	NYI	+17.5%
Torey Krug	BOS	+17.2%
Derek Ryan	CAR	+17.1%
Brayden Point	T.B	+17.1%

Worst 5v5 Net Threat Performances, 2017-2018

(thousand minute minimum)

Player	Team	Isolated Threat Relative to league average
Haydn Fleury	CAR	-27.5%
Justin Braun	S.J	-23.1%
Brooks Orpik	WSH	-19.3%
Dion Phaneuf	OTT & L.A	-18.9%
Viktor Arvidsson	NSH	-18.5%
Mattias Janmark	DAL	-17.7%
Mike Green	DET	-17.6%
Michael Del Zotto	VAN	-17.5%
Vladimir Sobotka	STL	-17.1%
Jonathan Drouin	MTL	-16.6%

Most Painful 5v5 Minutes, 2017-2018

(eight-hundred minute minimum)

Player	Team	*Context* Threat Relative to league average
Brandon Sutter	VAN	-27.7%
Devante Smith-Pelly	WSH	-23.1%
Carl Soderberg	COL	-22.9%
Jean-Gabriel Pageau	OTT	-22.8%
Marc-Edouard Vlasic	S.J	-22.2%

Most Sheltered 5v5 Minutes, 2017-2018

(eight-hundred minute minimum)

Player	Team	*Context* Threat Relative to league average
Brayden Schenn	STL	+25.3%
Mikhail Sergachev	T.B	+22.1%
Tyler Bozak	TOR	+22.0%
Jake Guentzel	PIT	+21.4%
Tomas Plekanec	MTL & TOR	+20.8%

Future Work

For shot density isolation itself:

- Non-linear effects. (Chemistry!)
- More subtle priors (including joint priors)

For a broader evaluation scheme:

- Special Teams (same machinery should work!)
- Goalies and shooting talent (totally different)

Thanks!