Dynamic Matching Pennies with Asymmetries: An Application to NFL Play-Calling

Matthew Mitchell*

March 17, 2010

Abstract

I introduce a model of matching pennies with dynamics and asymmetries across the actions. I show that returns may not be equalized across actions and that strategies may generate serially correlated actions. This is consistent with recent evidence from the NFL documented in Kovash and Levitt (2009)

1 Introduction

In a provocative recent paper, Kovash and Levitt [2] (KL) argue that NFL teams do not play optimal strategies. Their evidence is primarily two fold. First, the returns to running and passing plays is not equalized, as is the standard prediction in a game where players mix between strategies in a game. Second, offensive play calling is negatively serially correlated, suggesting exploitable patterns for the defense.

The goal of this note is to describe a two simple modifications to the author's game structure, each of which is results in a model that can match the qualitative facts from the NFL. Both add a dynamic element. The first example relies on an investment component to running; running early makes later running plays more effective, as if running "wears down the defense." The

*Rotman School of Management, University of Toronto. I thank Will Strange for helpful comments.

second example relies on asymmetric information, and embodies a popular football idiom: that you "run to set up the pass."[3]

KL use a matching pennies structure to explain the prediction of equal returns across plays. Suppose two players repeat the following game:

$$\begin{array}{ccc} & R & P \\ R & 0,0 & H,-H \\ P & L,-L & 0,0 \end{array}$$

The interpretation is that the column player (the defense) can stop the offense if it chooses the defense to match the offensive play. If the strategies do not match, the offense gets a positive payoff (and the defense, as this is a zero-sum game, loses the identical amount), so H and L are both assumed to be positive. Here I allow $H \neq L$ so that the offense may prefer "success" with one play versus another, to stand in for differences across teams in proficiency in various aspects of the game; this is immaterial to the point in KL, but will be relevant in the second model I introduce. Of course the only equilibrium has a fully mixed strategy, and the expected payoff to the two actions must be equalized, at odds with the data.

The first model makes only a small change to the environment, assuming that runs have a dynamic consequence on the defense, making future runs more effective. Football is a violent, physical game, and running plays are particularly physical and violent. One goal of those plays is to stun the defense in both the short and the long run (quote**). I therefore augment the model so that runs make future successful runs especially successful. As a result, teams run in the first period in order to generate second period rewards, making first period runs appear to have a low return when one measures only the return to the play itself in terms of net change in field position. Such a measure doesn't account for this dynamic benefit of runs. Further, on second plays that follow runs, the special potency of runs makes the equilibrium mixture shift toward more *passes*.

The first model highlights a point about matching pennies games that seems at odds with football. If a team has greater proficiency at passing, even in a one shot game as above with $H > L$, then that team passes more in equilibrium. The reason is simple: if it did not, the defense would always defend the run. This is a strong prediction about full information matching pennies. One natural way to break that link is by introducing asymmetric information.

The second model I consider does exactly that. In particular, I introduce asymmetric information about the strength of the opponent. Of course much of the difference across teams is not private; however, some asymmetric information is consistent with the production process in the NFL. During the week, coaches work extremely long hours evaluating the opposition and trying to determine what plays will be most effective[1]. It is natural to assume that the study uncovers some valuable information, that is necessarily private.[2] I show that asymmetric information about an offense's strength can generate predictions consistent with the KL evidence.

The goal of this paper is not to argue that it describes football in any sort of detail. The motivation is simply to show that the facts reported in KL can very readily be explained by optimizing agents, with additions that bear some resemblance to football. A closely related paper to KL is the study of penalty kicks by Chiappori, Levitt and Groseclose [1]. In a similar example from sports that fits a "matching pennies" intuition, they find that players taking penalty kicks can not be statistically differentiated from the optimal strategy. This conforms to the natural intuition that penalty kicks have nothing like the a strategy that wears down the goalie, or any asymmetric information of the sort that football coaches dig up during their week of film study. As a result, the complete information matching pennies example fits the penalty kick example quite well. It may not fit football quite as well.

The examples developed here in a sense embody a more general principle that average and marginal returns differ.[3] Virtually every model of optimizing behavior that you could write down would have the prediction that marginal returns across activities should be equalized. However, it may be the case that marginal returns are not indicative of average returns. The usual story is decreasing returns, where average returns overstate marginal returns. In these dynamic models, since strategies link actions across time, sometimes linking early runs with later passes, it isn't necessarily the case that a measured average returns are an exploitable margin for the player. Changing one play from run to pass in isolation does not change the payoff by the difference between the average return across plays.

[1] NFL coaches are noted for long hours in season. [4]

[2] For Super Bowl XLIV, the New Orleans coaching staff had, supposedly, installed different defensive plans for the first half, third quarter, and fourth quarter, in an attempt to attack different weaknessess they had uncovered. [5]

[3] Probably the best example of this difference is the marginal Q/average Q example.

2 Example 1: Running to Wear Down the Defense

One adage in football is that you run to wear down the defense. Here I model that as running plays making the defense more vulnerable to further physical running plays. Specifically, suppose the game is played twice. In the first period the game is as described above, with symmetric $H = L = 1$ for simplicity:[4]

	R	P
R	$0, 0$	$1, -1$
P	$1, -1$	$0, 0$

We denote the offense's strategy a "playcall" and the defenses strategy a "defense." The game is the same in the second play when following a pass. After a run, however, running becomes more efficient for the offense when successful, as the defense is "worn down":

	R	P
R	$0, 0$	$1 + x, -1 - x$
P	$1, -1$	$0, 0$

where $x > 0$.

2.1 Second Play

In the second period, after a run, the defense and offense both play mixed strategies. Denote by o the offense's probability of a pass, and d the defenses probability of defending the pass.

To compute the probability of the offense choosing a pass, we need to look at the defense's indifference condition, as both players will choose a mixed strategy. After a pass, indifference for the defense requires that

$$o = (1 - o)$$

so $o = 1/2$. Similarly, after a run, the condition is

$$o = (1 - o)(1 + x)$$

[4] Everything generalizes in the H and L example, at only a cost of greater notation and longer formulae.

chance of passing is $o = (1+x)/(2+x) > 1/2$, so passes are more likely after a run. *There is negative serial correlation of play calling from runs in the first period.*

The return to plays are identical across run and pass for each history. Returns differ across histories, however; to compute these returns, first compute that the defense's probability after a pass, computed from the offense's indifference condition, satisfies

$$d = (1 - d)$$

so $d = 1/2$. After a run the defense defends the pass with a lower probability $1/(2+x)$. The offense gets $1/2$ on plays that follow passes, and $(1+x)/(2+x)$ on plays that follow runs. The extra returns following runs is denoted V, i.e.

$$\begin{aligned} V &= \frac{1+x}{2+x} - \frac{1}{2} \\ &= \frac{x}{4+2x} > 0 \end{aligned}$$

One the second play, therefore, *passes are more prolific than runs*, since they are relatively more likely to occur in the state (after a run on the first play) where plays are more productive overall. Note that if one conditioned on the full history, the payoffs should be identical; however, in a many period model this would require keeping track of much more than just down, distance, and field position.

2.2 First Play

On the first play, players realize there is an extra benefit V accruing on the second play if the offense chooses a running play on the first play. Therefore the total payoffs can be described as:

	R	P
R	$V, -V$	$1+V, -1-V$
P	$1, -1$	$0, 0$

Since $V < 1$, once again there is only a completely mixed equilibrium. The defense's indifference condition is

$$o + (1-o)V = (1-o)(1+V)$$

So the relative frequencies of run and pass are the same as on the second play following a pass, $o = 1/2$.The offense's indifference condition is

$$d(1+V) + (1-d)V = (1-d)L$$

so $d = (1-V)/2 = (x+4)/(4x+8)$.

In terms of first period payoffs (that is, not including the returns V that accrue in the second period), runs make $(x+4)/(4x+8)$ and passes make $(1+V)/2 = \frac{4+3x}{4x+8}$. Note that once again passes outgain runs, conditional on being the first play. Since runs are the "unsuccessful" first play and also the driver of negative serial correlation, we see the feature reported in KL: *negative serial correlation is associated with unsuccessful first plays.*

One might think that indifference between plays implies that, in period one, the total returns across the two periods must be equalized across plays. This is not true because some of the reward to runs in the first period accrues as successful passes in period two. As a result, unconditionally across the two plays, *runs are outgained by passes.*

3 Example 2: Running to Set Up the Pass

The preceding example relies on the notion that running wears down the defense, which makes runs more productive. However, another common idea is that runs "keep the defense honest," so that later passing plays will be more effective due to their unpredictability. Here we model that not as a change in the fundamentals, but rather in the form of a two play game where the payoffs are unknown to the defense.

There are two types of offenses. The first, termed the "run type" (and denoted r) plays the game

	R	P
R	$0,0$	$H,-H$
P	$L,-L$	$0,0$

On the other hand, if the offense is "pass type" (denoted p):

	R	P
R	$0,0$	$L,-L$
P	$H,-H$	$0,0$

The game is repeated twice with whichever payoffs correspond to the offense's type. It is not critical that the two plays have the same payoff (H) for the favorable play and (L) for the unfavorable play; all that is material is that different teams have different strengths. The defense does not know of what the type is the offense. For simplicity I assume that the defense does not observe the payoff in the first period, so that it does not need to update based on the success of the first play; it uses only the play call in updating its beliefs μ about the chance that the offense is a pass type. This assumption can be easily relaxed in a model where the payoff was not a complete "giveaway" about the offenses type. Here the assumption stands in for the fact that the outcome of the play is very noisy relative to the information contained in the play call itself, which is observable. Such noise could easily be added, leaving the payoffs as the expected outcomes. We refer to a play of R by type r, or of P by type p, as *playing to strength.* We denote the initial beliefs of the defense that the offense is type r by μ, and their beliefs after observing the first play by μ'.

3.1 Stage 2 (last stage)

Here we compute the equilibrium for the situation where a defense with beliefs μ' finds itself in the final period. The offense of type i plays P with probability o_i, the defense plays P with probability d. At most one type can possibly mix, since if d is such that one type of offense is indifferent, the other type plays to strength. The defense must mix in any equilibrium, since, as usual, pure strategy plays by the defense are always exploited by the offense, making them suboptimal for the defense. For the defense to find it optimal to mix, there are two possibilities. The first is that $\mu' = 1/2$. The second is that at least one type mixes and $\mu' \neq 1/2$. We study that case first. The equilibrium we construct has the predominant type (the one that the defense believes to be more likely) and defense mixing, with the "rare" type playing its strength always.

Without loss we focus on $\mu' > 1/2$, so that run types are predominant. For the run offense to be indifferent, it must be the case that

$$dH = (1 - d)L$$

and therefore

$$d = \frac{L}{L + H} < 1/2$$

For the defense, indifference between defending run and pass implies (for simplicity let $o_r = o$, as only the run type is mixing)

$$\mu' oL + (1-\mu)H = \mu'(1-o)H$$

The left hand side is the payoff to playing run defense and the right is for pass defense. This simplifies to

$$o = \frac{H}{L+H}(2 - 1/\mu')$$

Note that this is between zero and one since $\mu' > 1/2$; $\mu' < 1/2$ makes indifference impossible, implying that one cannot construct an equilibrium where the rare type mixes and the predominant type plays to strength, as is intuitive. This implies that this equilibrium is unique for $\mu' \neq 1/2$: the defense must mix (or else everybody plays the opposite, and he should switch), and both offense types cannot mix simultaneously, nor can it be that the "rare" type is the one who mixes.

Here the run offense (which is typical) gets payoff $\frac{LH}{L+H}$ (since, for instance, if it runs is gets H with probability d, and both strategies yield the same payoff) and the pass offense gets $\frac{HH}{L+H}$, since its passes succeed $1-d$ of the time, each making H. Note that the rare type gets a higher payoff in equilibrium.

For $\mu' = 1/2$, any d between $L/(L+H)$ and $H/(L+H)$ is an equilibrium. Both types play to strength, and the defense is indifferent because of the equal shares of the two types. Note that for the extreme d in the range the type that is likely to be defended is just indifferent between strategies, but cannot mix in equilibrium and maintain mixing by the defense. There cannot by equilibria with d outside the range, since then both types play the same strategy and the best response for the defense is no longer to mix.

The key feature is that, for $\mu' = 1/2$, there are a continuum of equilibria that deliver different payoffs to the different types of offenses (but the same payoff, in expectation with respect to beliefs μ', for the defense). The payoff for the running team is dH, and for the passing team is $(1-d)H$. Note that the total payoff for the two teams is always H. We will exploit this property when constructing stage 1 equilibria, as the payoff in stage 2 can be chosen from a range of payoffs if the first period actions lead to $\mu' = 1/2$. We will then use the particular payoffs in stage 2 to support those actions as optimal in stage 1.

3.2 Stage 1

3.2.1 $\mu = 1/2$

A seemingly simple case is $\mu = 1/2$. In that case one might think there is a separating equilibrium where each offense plays its strength on play one, revealing its type, and then mixes, playing its weakness with probability $\frac{H}{L+H}$ in the second period (which its strength is defended with that probability). There is not, however, and the construction both helps explain the equilibrium that we construct, and rules out another possible type of equilibrium in the process. The payoff for each type in the proposed equilibrium is

$$\frac{1}{2}H + \frac{LH}{L+H}$$

However, the offense does not want to reveal its type in the first period. Instead it wants to "mimic" in the first period; the payoff from mimicking (playing its weakness, and then playing to strength in the second period) is

$$\frac{1}{2}L + \frac{HH}{L+H}$$

Which has difference

$$\begin{aligned}
&\frac{1}{2}H + \frac{LH}{L+H} - \frac{1}{2}L - \frac{HH}{L+H} \\
&\frac{1}{2}(H-L) - \frac{H}{L+H}(H-L) \\
&(H-L)(\frac{1}{2} - \frac{H}{L+H}) < 0
\end{aligned}$$

An equilibrium, instead, is one where both types mix with identical probabilities in the first period, and the defense mixes with $d = 0$.[5] In the second period, since μ' is still 1/2, there are many possible payoffs to the one shot game. The equilibrium chooses second period payoffs so that mixing is optimal in the first period; both types play to strength in period 2, but receive different payoffs from the set of equilibrium payoffs described above for the second stage when $\mu' = 1/2$..

[5] Note that whenever both types play the same probabilities, only $d = 1/2$ can make the two types indifferent simultaneously.

Mixing by each team in the first period requires

$$\frac{1}{2}H + V_H = \frac{1}{2}L + V_L$$

Where V_H is what a team gets in the second period after playing its strength in the first period, and V_L is what the team gets in the second period after playing weakness in the first period.

In the second period, the total payoffs must add up to H, i.e.

$$V_H + V_L = H$$

So $V_H = (H + L)/4$ and $V_L = (3H - L)/4$. It is easy to verify that $H/(H + L) > V_L > V_H > L/(H + L)$, so the payoffs in the second period that are required are feasible equilibrium payoffs in the second period when $\mu' = 1/2$. Note that the defense does the opposite of the first period offenses play call in the second period with probability $(1 - L/H)/4 < 1/2$, so play calling by the defense is positively correlated with last period's offensive call.

3.2.2 Intermediate μ

Suppose that μ is greater than $1/2$, but smaller than $\frac{2H+L}{3H+L}$. We construct an equilibrium that has both types mixing in period one, but now with different probabilities so that $\mu' = 1/2$ after a pass (the unlikely strength) and $\mu' > \mu > 1/2$ after a run. In the latter case, the second stage can proceed in only one way: the running team mixes, and the unlikely passing team always passes. The payoffs are $\frac{HL}{H+L}$ for the running team and $\frac{HH}{H+L}$ for the passing team in the second stage in that case.

In the case where $\mu' = 1/2$ in the second period, the payoffs are chosen from the set of feasible payoffs $[HL/(H + L), HH/(H + L)]$ so that mixing is optimal in period one for the type that is likely. Let V be what the ex ante predominant type (the running team) gets in that situation (with the other type getting $H - V$, since any period two equilibrium with $\mu' = 1/2$ has total payout H across the types). In order for mixing by each type to be optimal, it must be the case that

$$\begin{aligned} dH + \frac{HL}{H+L} &= (1-d)L + V \\ dL + \frac{HH}{H+L} &= (1-d)H + (H - V) \end{aligned}$$

where the first condition is indifference for the run (predominant) type and the second is for the pass type. An immediate implication is that $d = 1/2$ and

$$V = \frac{1}{2}(H - L) + \frac{HL}{H + L}$$

Note that, with $d = 1/2$, in this equilibrium the worst first plays (in terms of play one returns) are plays to weakness; these are always followed by plays to strength. In other words, *there is strong negative serial correlation for plays that turn out badly in the first period.*

To get a sense of how different plays compare, note that $V > H - V$ implies that the predominant type does well in period two; since everyone plays to strength, this implies that those plays do well in period two. Once again, there is *no equivalence between running and passing plays*, even conditional on the play number (which one might think of as akin to conditioning on down and distance in KL). We discuss the relative performance of the various types of plays below.

To close this model, and to be able to calculate the returns to various sorts of plays, we need to know the mixing probabilities in period one. There are two equations that determine these. The first is the indifference condition for the defense:

$$\begin{aligned} & \mu \left(o_R L + o_R V + (1 - o_R)\frac{HL}{H + L} \right) + (1 - \mu) \left(o_p H + o_p(H - V) + (1 - o_p)\frac{HH}{H + L} \right) \\ = \ & \mu \left((1 - o_R)H + (1 - o_R)\frac{HL}{H + L} + o_R V \right) + (1 - \mu) \Bigg((1 - o_p)L + (1 - o_p)\frac{HH}{H + L} + o_p(H - V \end{aligned}$$

or

$$\mu o_R L + (1 - \mu) o_p H = \mu(1 - o_R)H + (1 - \mu)(1 - o_p)L$$

The second condition is the one that says the optimal beliefs after a pass are $\mu' = 1/2$. Applying Bayes' Rule,

$$\frac{\mu o_R}{\mu o_R + (1 - \mu) o_p} = \frac{1}{2}$$

so

$$(1 - \mu) o_p = \mu o_R$$

Replacing in the indifference condition:

$$
\begin{aligned}
\mu o_R L + \mu o_R H &= \mu(1-o_R)H + (1-\mu)L - \mu o_R L \\
2\mu o_R(L+H) &= \mu H + (1-\mu)L \\
o_R &= \frac{1}{2}\frac{H + \frac{1-\mu}{\mu}L}{(H+L)} \\
o_p &= \frac{1}{2}\frac{\frac{\mu}{1-\mu}H + L}{(H+L)}
\end{aligned}
$$

Since $\mu > 1/2$, $o_R < 1/2$ and $o_P > 1/2$; running teams tend to run in the first period, and passing teams tend to pass. Here serial correlation is, overall, positive: runs in the first period indicate a running team, who is more likely to run in the second period (since passing teams always pass in the second period).

Next we turn to payoffs. In the first period, since $d = 1/2$, every team gets $H/2$ for plays to strength, and $L/2$ for plays to weakness. The return per run is

$$\frac{\mu(1-o_R)H/2 + (1-\mu)(1-o_P)L/2}{\mu(1-o_R) + (1-\mu)(1-o_P)}$$

and the return per pass is

$$\frac{\mu o_R L/2 + (1-\mu)o_P H/2}{\mu o_R + (1-\mu)o_P}$$

Both are convex combinations of $H/2$ and $L/2$; the return per pass puts weight

$$\frac{(1-\mu)o_P}{\mu o_R + (1-\mu)o_P}$$

one $H/2$. But that is $1/2$ according to the values of o_i above. The return per run puts weight

$$\frac{\mu(1-o_R)}{\mu(1-o_R) + (1-\mu)(1-o_P)}$$

on the $H/2$ term; it is easy to verify that this is greater than $1/2$. This implies that *runs outgain passes on play one.*

Note, as mentioned, following a pass, the payoff to running teams in period $2(V)$ is higher than the return for passing teams in period 2 $(H-V)$.

Since both teams play to strength following a pass, and following a pass teams are in equal proportions (as $\mu' = 1/2$), this implies that *runs make higher returns following a pass.*

On the other hand, following a run, passing teams always pass and get a payoff $HH/(H+L)$; running teams get $HL/(H+L)$ for either of the strategies between which they mix. Therefore passes that follow runs include both high $HH/(H+L)$ and low $HL/(H+L)$ payoff plays, but runs that follow runs are all by the running team and earn $HL/(H+L)$, and therefore *passes outgain runs after a period one run.*

To see how this can mean that passes dominate runs in period two unconditionally, suppose L is nearly zero, and μ is just slightly above 1/2. In this case, the payoff in the state following a pass is essentially the same to run and pass types (since $L \approx 0$, the difference between V and $H - V$ is small); however, after a run, the pass team makes $HL/(H+L)$ (all from passes) and the running team makes $HL/(H+L)$ (approximately zero). It is easy to confirm that passes outgain runs in period 2 in that case. To put it in terms of the conditional returns to runs and passes, when μ is near 1/2, the two conditions (period one run and period one pass) are nearly equally likely, and all that matters is the relative payoffs to the plays in the different states. After a pass, runs and passes make almost identical returns when L is small; however, the difference between passes and runs following running plays remains substantial (approximately H). Further, in this example, the return to runs and passes is virtually identical in the first period, since both are played almost the same amount; therefore across all plays passes outgain runs.

If, contrary to the example with L small and μ near 1/2, runs were to outgain passes unconditionally in period 2, then there is still a sense in which the model would be on track to match the facts in KL: the model just needs to switch the "identities" of run and pass, so that the pass type is predominant, and the rewards to the passes would dominate runs. Therefore the payoffs between runs and passes might inform us as to whether teams are running to set up the pass, or passing to set up the run.

3.2.3 Extreme μ

When μ is sufficiently large, we can construct equilibria where the defense has $\mu' > 1/2$ after both run and pass. The occurs because the initial beliefs μ are so extreme that beliefs $\mu' = 1/2$ are impossible when the predominant (running) team mixes even if the pass team does not. The indifference condition for the defense in this case is (where, again, we have used o for the running team's mixture since the passing team does not mix):

$$\mu\left(oL + \frac{LH}{H+L}\right) + (1-\mu)\left(H + \frac{HH}{H+L}\right) = \mu\left((1-o)H + \frac{LH}{H+L}\right) + (1-\mu)\frac{HH}{H+L}$$

Which reduces to the second period condition,

$$o = \frac{H}{L+H}(2 - 1/\mu)$$

The mixing condition for the run offense is

$$dH + \frac{LH}{H+L} = (1-d)L + \frac{LH}{H+L}$$

so again

$$d = \frac{L}{L+H}$$

Now the final question is what μ would lead to $\mu' > 1/2$. Clearly runs lead to $\mu' = 1$, since only running teams run. For passes the condition is

$$\frac{\mu o}{\mu o + 1} > 1/2$$

or

$$\mu > \frac{2H+L}{3H+L}$$

As a result, this sort of equilibrium cannot occur for the range of μ considered in the previous subsection. Note that in this equilibrium play calling is not serially correlated but unlikely strengths have higher payoffs (passing if $\mu > 1/2$, and vice versa). All plays for the running team yield $\frac{LH}{L+H}$ (which are a mix of runs and passes), whereas the passing team gets $\frac{HH}{L+H}$ from all of its plays (which are all pass plays). Therefore in this example *passes outgain runs.*

3.3 Discussion

It may seem counterintuitive that one play could have a higher return. But recall the first period indifference condition. It says that the total payoff to each play (summed across plays) is equal between runs and passes. However, the return to period one runs is partially capitalized in period two passes, and vice versa. Therefore, when one measures the return from a given play, he is also measuring the capitalized payoffs to earlier plays that set up the later play, by impacting beliefs. This is exactly the idea behind "running to set up the pass."

Notice that, in this example, which play dominates depends on situation as well as parameters like μ. The relative value of the plays can shift dramatically, for instance, as we move from the intermediate to the extreme μ examples. Further, negative serial correlation can emerge for bad plays, while not emerging overall in some cases. The richness of the model suggests it can generate interesting implications that might be testable. Foremost is the "surprise" aspects: does changing play calls yield higher payoffs? Here, changing playcalls generally signals that the first play was hiding the true strength of the team.

4 Summary

By adding a dynamic element to the matching pennies environment, one can model features like investment and learning that may be relevant for empirical evidence. I show here that such changes can explain the empirical results of Kovash and Levitt (2009) from the National Football League in a way that conforms to at least some rough intuition about football Interestingly, the two modifications have some predictions that differ; for instance, the "run to set up the pass" model has strong negative correlation for unsuccessful plays, but actually predicts positive correlation overall, counter to the facts. These differences would allow the two models to be put together, and the relative importance of each force to be estimated. A more detailed model of football, including the possibility that the data in KL might inform such a model and even allow for estimation of a model, is left for future work.

References

[1] Chiappori, P.A., S. Levitt and T. Groseclose (2002). "Testing Mixed-Strategy Equilibria When Players Are Heterogeneous: The Case of

Penalty Kicks in Soccer," American Economic Review 92(4) pp. 1138-1151.

[2] Kovash, Kenneth and Steven D. Levitt (2009). "Professionals do not Play Minimax: Evidence from Major League Baseball and the National Football League," NBER Working Paper 15347.

[3] *Offensive Philosophy (American Football),* http://en.wikipedia.org/wiki/Offensive_philosophy_(American_football)).

[4] Peters, Justin (2006) "No Sleep Till Touchdown: NFL coaches, the hardest-working men in human history" http://www.slate.com/id/2149181/

[5] Transcript of New Orleans Saints postgame press conference, availaible at http://www.nola.com/saints/index.ssf/2010/02/new_orleans_saints_defensive_c_14.html