

PROOF THAT NICK "STOXTRADER" GRUDZIEN, "LITTLEZEN," AND "KINETICA" COLLUDED

NOAH STEPHENS-DAVIDOWITZ

ABSTRACT. I show to an extremely high level of statistical significance that Nick "Stoxtrader" Grudzien played substantially differently than LittleZen and Kinetica and vice versa. The way in which they played against each other is consistent with softplaying and completely inconsistent with other plausible explanations. In addition, their play against each other was normal in situations in which cheating was not possible (i.e. when no other players were in the hand). I conclude that they colluded.

This contradicts the results of investigations of both PokerStars' and Full Tilt Poker's security departments, both of which found no evidence of cheating in spite of their access to much more data than I had.

1. HISTORY

Nick "Stoxtrader" Grudzien was recently accused by an anonymous source of multiaccounting, using the accounts stoxtrader, 40putts and bulltf0rdtuff on FTP and stoxtrader, knockstiff, and gr3atvlewbr0 on Stars (see this long and complicated thread: <http://forumserver.twoplustwo.com/19/high-stakes-pl-nl/stoxtrader-cheating-multi-accounting-discussion/>). He admitted to that (<http://forumserver.twoplustwo.com/showpost.php?p=17566569/&postcount=971>).

However, in March of 2009, the accounts 40putts and knockstiff had been accused of colluding with Kinetica and LittleZen respectively by a few different HSNL regs (see posts scattered throughout this thread: <http://forumserver.twoplustwo.com/19/high-stakes-pl-nl/hsnl-march-2009-thread-424702/index16.html>). Both Stars and FTP investigated and

found them innocent, but the fact that 40putts and knockstiff were multiaccounts made people more suspicious, so the community started looking into it more.

People found some anecdotal evidence that the accused accounts were softplaying. In other words, they thought they were intentionally avoiding taking money from each other in order to take more money from others at the table. Some early statistical evidence suggested that this might be true, but the sample sizes were way too small. From what I gathered from the thread, most people assumed he was guilty.

This situation really bothered me. There were only two possibilities. Either both Stars AND FTP security had failed to identify collusion that had been pointed out to them by numerous reputable people, or the court of public opinion had convicted an innocent man.

The first case is a big problem for obvious reasons. The second might not seem so bad, since stoxtrader had already admitted to a lesser crime that completely ruined his reputation, and he probably wasn't really being punished any more just because everyone thought he'd done something else wrong. But, if Stars and FTP were correct but nobody believed them, it would unjustly cost them credibility, which would be terrible for the poker community at large. In short, it's bad for everyone when the players and the sites disagree.

So, I wanted to resolve this. Even more importantly, I wanted to prove that such issues could be resolved. Without such proof, there's nothing to hold the sites accountable, and no reason for the public to trust any facts that they present. The following is my resolution of the case at hand and proof that these things can be resolved.

2. DISCLAIMER

In this post, I'm making public a lot of data about how other players play. I realize that there are ethical considerations here, and I want to address them. The data that I'm posting about 40putts, Kinetica, knockstiff, and LittleZen is pretty extensive and would likely be useful to their opponents. Stoxtrader has admitted that he multiaccounted with the accounts 40putts and knockstiff, so as far as I'm concerned, he's forfeited the right to

privacy there. I'm posting the data about LittleZen and Kinetica because it shows that they cheated, so they've forfeited their right to privacy as well.

In order to make my case against those four accounts, I have to show that they played differently against each other than against other similar players. In order to do that, I have to present a small amount of information about other players that they played against. I've done my best to present as little of this data as I felt was necessary and to present it in such a way that makes it difficult to use for other purposes. I'm sure that if someone really wanted to, it wouldn't be too hard for them to figure out the PFR and fold to 3-bet of some of these players, and I apologize for that. If anyone would like their information removed from here, please contact me, and I will remove it. But, of course once I've made information public, I don't have complete control over it.

PTR has also agreed not to release any of the data they gathered for me without my consent.

3. DATA

Except where I say otherwise, the data that I'm presenting comes from queries that Poker Table Ratings ran on their database, which has 500,547 hands on LittleZen, 183,608 hands on knockstiff, 265,484 CAP hands on Kinetica (634255 total hands), and 222,683 CAP hands on 40putts (425259 total hands). I got the data without screen names (except for the four suspect accounts) and do not have the hand histories that it came from. I do have my own large database of hands that I gathered to look into this, and their data agrees very well with the data that I have. I've also looked at some subsets of the database provided by PTR and grilled them pretty thoroughly about their collection methods. So, I believe PTR's data to be true and unbiased.¹

¹In the interest of full disclosure, I should say that I originally thought that the two databases did not agree. It turns out that there are a number of biases in my database that are not present in the PTR database. My database has much fewer 50 BB+ tables than PTR's, fewer short-handed tables, and the coverage varies over time differently than PTR's coverage. I can account for all of these biases because of the way that the data was gathered, and I believe them all to be biases on my end and not on PTR's. After

All data comes from 5/10 to 25/50 games with between 4 and 6 players at the table. Except where I say otherwise, the confidence intervals in the graphs use the normal approximation to the binomial distribution and represent 1 standard deviation, so $\approx 68.2\%$ confidence intervals. This is standard practice in cases where $n * p > 5$. My original forum post contains a more nuanced discussion of the statistics involved.

Any individual calculation (i.e., when I say "the odds of this happening randomly are one in a million") was done using exact binomial distribution calculations. So they are all perfectly accurate except for some rounding.

4. LITTLEZEN

LittleZen is a professional shortstacker on Stars. Of the 500,547 hands that PTR has on him, 323,325 (64.6 %) were played with less than 25 BBs. Of the 30,180 hands that PTR has on both him and knockstiff, LittleZen had less than 25 BBs in 19,186 (63.6 %). In my own database (which is missing a lot of hands at 50 BB+ tables), I have 529,531 hands on him. 463,976 (87.6 %) of those were played with less than 50 BBs, 430,898 (81.4 %) were played with less than 30 BBs, and 393,266 (74.3 %) were played with less than 25 BBs.

The fact that LittleZen is a professional shortstacker is very important for my case. If he were playing with 100 BBs, then he could have some sophisticated reasoning for playing very differently against different players' open raises. But, with less than 25 BBs or so, a reasonable player almost always just shoves or folds against an open raise. In my database, LittleZen's response when someone open raised when it folded to him was to shove 10 % of the time and call 0.8 % of the time, so he clearly agrees. Whether or not a good player shoves in a spot like this depends on a simple equation based on how often he expects his opponent to fold, what range he expects to call him, and how many players are left to act.


As a professional mid/high-stakes shortstacker with a decent winrate (1.54 PTBB/100 or

rigorously accounting for these biases, the two databases appear to agree very well where I tested them. PTR had no prior knowledge of what data I had or what I would look for to confirm accuracy. If people are curious, I'll happily show my work.

3.08 bb/100), LittleZen certainly knows this basic concept. And, he applies it very well and very consistently, except against knockstiff (stoxtrader's Stars multiaccount).

Figure 1 shows LittleZen's 3-bet % from positions other than the BB with no callers (i.e. not a squeeze) vs. his 49 most common opponents' open raises (excluding opponents with VPIP > 30 %). Knockstiff is in red and is the first datapoint. The rest of the players are in order of decreasing sample size (i.e. opportunities to 3-bet). One player has a larger sample size than knockstiff. 60

FIGURE 1. LittleZen 3-bet % vs. Open Raises (all players, no callers, not BB)



Notes: You should be looking at the y-axis here. The x-axis is just used to separate out the data a bit.

The error bars represent about a 68.2 % confidence interval (calculated by $\text{SQRT}(3\text{-bet \%} * (1 - 3\text{-bet \%}) / \text{sample})$). Double the error bars and you get a 95.5 % confidence

interval. The lack of confidence comes from the fact that LittleZen might 3-bet someone with x % of hands on average, but he might have been dealt those hands more or less than x


Notice that LittleZen's 3-bet % is quite consistent for everyone except knockstiff. This is in line with the description I gave earlier of a professional short-stacker's play. Against knockstiff, he 3-bets only 2 % of hands on average. For reference, KK+/AK is already 2.1 % of hands.

Also notice how small the error bar is on knockstiff (± 0.36 %). This is because of the very large sample of hands that they've played with each other and the fact that error tends to be smaller with smaller 3-bet %. It's that small error that allows me to say confidently that LittleZen plays differently against knockstiff. Indeed, the next lowest 3-bet % is 4.4 %, or about 6.6 standard deviations from the mean. At this point, the normal distribution isn't a very good approximation, but luckily I can calculate the exact binomial approximations. The odds that LittleZen's actual 3-bet % against knockstiff is 4.4 % and he just happened to show up at 2 % in PTR's sample are extremely low (about one in 3.5 million).

The data is even more extreme if I compare with the (weighted) average of the 3-bet %s of the other 48 players. That number is 8.0 %, 16.6 standard deviations away from the mean. The odds that LittleZen's 3-bet % vs. knockstiff was actually that high are extremely low (about one in $663 * 10^{21}$, or 663 sextillion. In naming these numbers, I've used the American system. Sorry British people.)

So, maybe LittleZen plays differently against stoxtrader's account, knockstiff, for some legitimate reason. As I explained earlier, a shortstacker's play is very simple, so there aren't many variables. Figure 2 shows LittleZen's 3-bet % in the same situation against these same players, now with pre-flop raise % on the x-axis.

FIGURE 2. LittleZen 3-bet vs. PFR (All Players, not BB, no callers)




Knockstiff's preflop raise % is not extremely low or high. I have data on many players with very similar preflop raise %, but LittleZen played very differently against them. Figure 3 shows the same data broken up by fold to 3-bet %.

Knockstiff's fold to 3-bet % is actually one of the highest in the sample. This should make LittleZen more likely to 3-bet him. You can even see a general trend in the graph. But, he actually 3- bets him much less than everyone else.²

²This could be due to knockstiff playing strangely against LittleZen's 3-bets, but that is not the case. Knockstiff folds to LittleZen's 3-bets 73.8 % of the time.


FIGURE 3. LittleZen 3-bet vs. Fold to 3-Bet (All Players, not BB, no callers)



So, maybe it's a combination of these two stats that causes him to 3-bet so infrequently. Figure 7 shows the same data filtered for only people with PFRs within 1 % of knockstiff's and fold to 3-bets within 10 %:

A combination of both stats doesn't explain his strange play either. The next lowest value in this group is 6.9 %. The odds that LittleZen actually 3-bet knockstiff 6.9 % of the time by PTR's data showed 2 % are extremely low (about one in $977 * 10^{15}$, or 977 quadrillion). If I take a weighted average of LittleZen's 3-bet % vs. these players (excluding

FIGURE 4. LittleZen 3-bet % vs. Opponents Similar to knockstiff (PFR within 1 % and F3 within 10 % of knockstiff, not BB, no callers)




knockstiff), I get 9.4 %. The odds that LittleZen actually 3-bet knockstiff 9.4 % of the time by PTR's data showed 2 % are extremely low (about one in $54 * 10^{30}$, or 54 nonillion).

So, the two primary variables that a shortstacker should use to decide his shoving range do not explain LittleZen's suspicious play against knockstiff. But, maybe LittleZen thinks differently than I do. If he did, you'd expect this difference to be reflected in LittleZen's big blind shoving range as well. However, if he's colluding with knockstiff, he would have no reason to 3-bet him less from the big blind since there are no players left to act in that spot. Figure 5 shows how he actually played in these situations. (Note that I've switched the axes.)

My sample on hands in the big blind is obviously not as big as hands in all other positions, so the error bars are fairly large here. Still, it is clear that LittleZen's 3-bet % in the big blind vs. knockstiff is in line with his 3-bet % vs. other players. There are a

FIGURE 5. LittleZen 3-bet BB vs. 3-bet not BB



lot of players that he plays similarly against in the big blind but very differently in other positions. Because of the large sample with knockstiff, the error bar on him is actually fairly small, so I can say with high confidence that LittleZen's play in the big blind against him is not particularly strange.³

There's one other possible legitimate explanation for this data. I said earlier that this type of play is indicative of cheating with short stacks. It's possible that what I'm finding

³The few scattered players with extremely low or extremely high values in this chart are to be expected due to the large error and the fact that 49 players are shown.

is simply caused by LittleZen playing differently against knockstiff with a deep stack. Since LittleZen has roughly the ratio of hands played short stacked vs. deep stacked with knockstiff as in the entire sample, this would still represent suspicious play against knockstiff, but it would allow for a lot more legitimate explanations. (For example, maybe LittleZen thinks knockstiff plays well postflop in 3-bet pots and therefore avoids 3-betting.)

But, PTR was nice enough to provide me with the hands that they have in which knockstiff raises and LittleZen is left to act. I imported these into Holdem Manager, filtered for hands in which it folded to LittleZen with less than 25 BBs in a position other than the BB, and found that he 3-bet 3.5 % (with an error of +/- 0.67 %) of his 742 opportunities to do so. It's not surprising that this number is higher than the 2 % over the total sample since any good short stacker will tell you that a short stack should 3-bet more with 20 BBs than with 30. However, this 3-bet % is still significantly lower than his 3-bet % vs. the other players, even though the other players' data includes deep stacked play (in which LittleZen 3-bets much less in general).

So, LittleZen's play vs. knockstiff was much different than his play against other players. I have put a lot of time and effort into looking for a legitimate reason why this might be true, and I've found none. I've asked others to do the same, and they've found none. I've also shown that his play was extremely abnormal in situations in which such play would benefit a cheater (not in the BB), but perfectly normal in the most similar situation in which such play wouldn't be advantageous to a cheater (in the BB). So, the only reasonable conclusion is that LittleZen intentionally avoided playing aggressively against stoxtrader's account, knockstiff, thus costing himself a lot of money in order to allow knockstiff to take even more money from other players at the table. This is a clear case of collusion.

5. KNOCKSTIFF

One could argue that LittleZen has no incentive to softplay knockstiff unless knockstiff softplays against him or knockstiff pays him to do so in some way. Both of those things

constitute collusion. So, in my opinion, it would be perfectly reasonable to assume that knockstiff is guilty of some form of collusion simply based on the LittleZen data. However, the case would certainly be a lot clearer if I could show that knockstiff did in fact softplay against LittleZen.

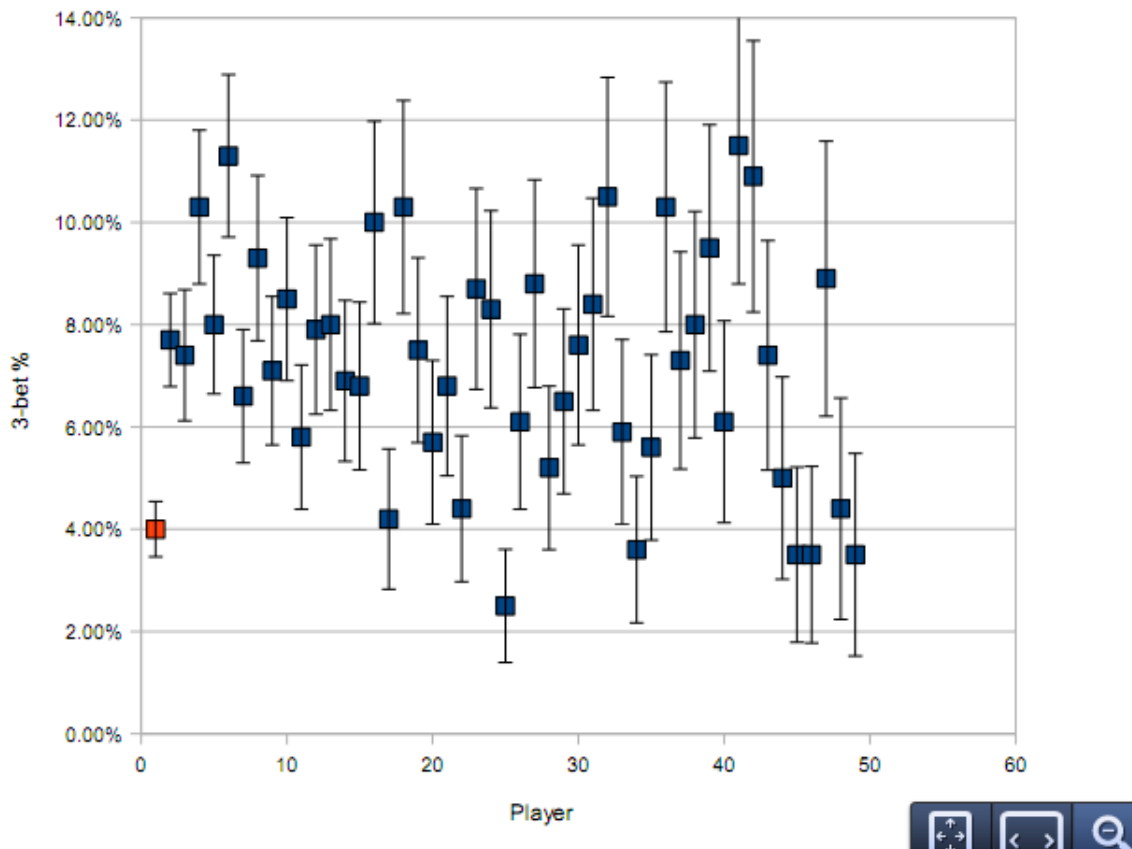
Unfortunately, the samples that PTR and I have on knockstiff are much smaller than the comparable samples on LittleZen. I have 234,460 hands and PTR has 183,608. Knockstiff also played with a deep stack more often than LittleZen, which further complicates matters. Of the 183,608 hands on knockstiff that PTR used for this investigation, 76,969 (41.9 %) were played with less than 25 BBs. Further complicating things, knockstiff varied his play by opponent a lot more than LittleZen did. Also, knockstiff was much less blatant about his collusion.

The result of all this is that my argument becomes more tedious. Figure 6 shows a graph of stoxtrader's account knockstiff's play against his 49 most common opponents (excluding players with ≥ 30 % VPIP).

knockstiff 3-bet % vs. Open Raises (All players, no callers, not BB)

As you can see, the error bars are quite large. They only represent a 68.2 % confidence interval, so the data is quick murky. Since I have 48 players other than LittleZen listed, it's not surprising that there are a few players who are one or two standard deviations below the mean. However, the fact that there are 8 players with a 3-bet % less than 5 % is a bit surprising at first. The weighted average of his 3-bet % against this sample (excluding LittleZen) is 7.53 %. If knockstiff didn't vary his play a bit by opponent (It seems like LittleZen barely did) and 3-bet each opponent roughly 7.53 % of the time in these spots, these numbers would be pretty unlikely. Figure ?? is a representation of that with the 8 players other than LittleZen whom he 3-bet less than 5 % of the time in this spot.

FIGURE 6. Knockstiff 3-bet % vs. Open Raises (All players, no callers, not BB)



Of course, those probabilities aren't THAT low, but they're bigger outliers than you typically see in this sort of data. So, knockstiff probably did vary his play against different players much more than LittleZen did.

However, if we just use 6 % instead of 7.53 %, the numbers fall in line. (See Figure 8.)

In a sample of 48 players, it's not at all surprising to see one result that had a 1.2 % chance of happening and 3 results with between 10 % and 20 % likelihood. This also makes intuitive sense. If knockstiff 3-bets 7.53 % on average in this spot, it wouldn't be surprising at all for him to 3-bet some players (maybe tighter players, for example) 6 % of the time,

FIGURE 7

Sample	3-bet %	Chance of this 3-bet % or lower occurring randomly from a real 3-bet of 7.53%
199	2.50%	0.20%
86	3.50%	10.40%
115	3.50%	6.00%
113	3.50%	6.60%
169	3.60%	2.60%
214	4.20%	3.60%
206	4.40%	4.80%
90	4.40%	18.40%

FIGURE 8

Sample	3-bet %	Chance of this 3-bet % or lower occurring randomly from a real 3-bet of 6%
199	2.50%	1.20%
86	3.50%	23.50%
115	3.50%	17.40%
113	3.50%	18.50%
169	3.60%	11.40%
214	4.20%	16.80%
206	4.40%	20.40%
90	4.40%	36.60%

and then it wouldn't be surprising that a few of these would end up as the low points on my graph.⁵

The odds that knockstiff 3-bet LittleZen with 6 % of hands on average but PTR's sample showed 4 % are about 0.1 % So, the case here is not nearly as strong as that against LittleZen. It's possible that knockstiff varies his play wildly against different players, though the data does support a more reasonable assumption. It's also possible that knockstiff just happened to run bad when LittleZen open raised.

If I saw this data independently of the data on LittleZen, I would declare it suspicious but inconclusive, and I would likely stop at this point. However, given LittleZen's softplaying against knockstiff, this data becomes much more suspicious, and I think it deserves a closer look. You'll see that after that closer look, things get a lot clearer.


Figure 9 shows knockstiff's 3-bet % (in positions other than the BB with no callers in between) vs. the same sample by their preflop raise %.

Most of the very low 3-bet %s are against extremely tight players. In fact, all of the players that knockstiff 3-bets less than LittleZen have much lower preflop raises. You can see in figure 9 the general trend that higher preflop raises lead to higher 3-bet %s. This is of course natural. LittleZen does not follow this trend, however. You can see that he's not the only player who doesn't follow this trend, but the others all have much larger error bars (because of much lower sample sizes). Figure 10 shows the same thing restricted to players with PFRs within 1 % of LittleZen's:

The weighted average 3-bet % for this set of players is 8 %. The odds that knockstiff actually 3-bets LittleZen with 8 % of hands in these spots on average but randomly has 4 % in PTR's DB are extremely low (about one in 187 million). However, this isn't particularly

⁵Note that the difference between knockstiff's 7.53 % average and the 6 % that I'm allowing for here and LittleZen's 8.0 % average and his 3-bet % against knockstiff of 2 %. Normal variation in how a reasonable player responds to different players is a part of the game and of course totally fine and expected. When the difference is very large and there is no reasonable explanation for such a large difference, something is wrong.

FIGURE 9. Knockstiff's 3-bet % vs. PFR (All Players, not BB, no callers)




decisive considering that there is one player whom he 3-bet 4.2 % of the time. (The odds that this happened randomly are 2.5 %. I already argued earlier that the data supports this player being 3-bet about 6 % of the time, but his presence still certainly hurts my argument.)

Figure 11 shows the same data broken up by fold to 3-bet %.

Note that LittleZen's fold to 3-bet is quite high, which should encourage knockstiff to 3-bet him more often than the field, not less. Of course, this argument is again not incredibly

FIGURE 10. Knockstiff 3-bet % vs. PFR (Players with PFR between 18.6 % and 20.6 %, not BB, no callers)



strong because there are again players with much smaller samples with a similar fold to 3-bets whom knockstiff 3-bet even less.

However, when I filter this data for players with similar PFRs to LittleZen, it gets a lot clearer. See Figure 12.

The graph shows a clear trend. Knockstiff 3-bets very lightly vs. the player who's raising almost 20 % of hands and folding less than half the time. These are presumably 3-bets for value. Then the 3-bet % comes down sharply as he encounters people against whom you can't value bet 10 % of hands. Then the 3-bet % rises fairly quickly as his 3-bets start to gain equity from pure fold equity. The error on the individual data points is pretty large, but the trend is clear in spite of that. Note in particular that the one player against whom knockstiff played very tightly has a much lower fold to 3-bet than LittleZen, in addition to his lower VPIP. So knockstiff has a legitimate reason to be very tight against him.

LittleZen, however, completely violates this trend. He has the highest fold to 3-bet of the entire group but gets 3-bet least often.

FIGURE 11. Knockstiff 3-bet % vs. Fold to 3-bet (All players, not BB, no callers)





Figure 16 shows the same data, but only for players with PFRs within 1 % of LittleZen's and fold to 3-bets within 10 %.

Keep in mind that these are players who play very similarly to LittleZen, so while there are only 8 (not counting LZ), I think this is quite alarming. In fact, it's even more alarming than it looks because LittleZen has the highest fold to 3-bet of any player in this sample. His fold to 3-bet is a little over 7 % higher than the group average. So, knockstiff should be 3-betting him significantly more often than the other players in the sample. Instead, he 3-bets them 7.99 % of the time (weighted average) and LittleZen 4 % of the time.

The odds of knockstiff's actual 3-bet % against LittleZen being the lowest number in this sample (5.7 %) but PTR's data showing 4 % over 1,311 opportunities are about one


FIGURE 12. Knockstiff 3-bet % vs. Fold to 3-bet (Players w/ PFR between 18.6 % and 20.6 %, not BB, no callers)



in 240. So not too unlikely, though definitely a pretty surprising coincidence. But again, this is a sample of players who all play very similarly to LittleZen. The 5.7 % number is only over 210 hands, which has about a 13.6 % chance of happening randomly if the real % chance were the weighted average of 7.99 %. The other low number (6.1 %) only has a sample of 196 hands, which has a 20.6 % chance of happening randomly, and he also has the lowest fold to 3-bet of the sample, only making the cut by half a percent. So it's not too surprising that these two players out of 8 were 3-bet a little less.

It is very surprising that LittleZen was 3-bet so much less. The odds of PTR showing 4 % over the 1,311 opportunities that knockstiff had to 3-bet LittleZen when the actual

FIGURE 13. Knockstiff 3-bet % vs. Fold to 3-bet % (Players w/ PFR and F3 similar to LittleZen, not BB, no Callers)




number is 6.1 % (the second lowest datapoint) are about one in 1700. The odds that PTR showed 4 % when the actual number is 7.99 % are extremely low (about one in 174 million).

It's possible there's something about LittleZen's play that I didn't consider. So in order to check for this, I look at knockstiff's play in the big blind against LittleZen's opens. If there's something I'm missing, this difference should be reflected in knockstiff's big blind shoving range as well. However, if he's colluding, he would have no reason to 3-bet him less from the big blind since there are no players left to act in that spot.

Figure 14 shows this data.

FIGURE 14. Knockstiff 3-bet BB vs. 3-bet Non-BB (All players, no callers)



Notes: These error bars aren't strictly accurate.

knockstiff 3-bet LittleZen's open raise 11.2 % of the time when it folded to him in the big blind. The weighted average for this statistic over the entire sample (excluding LittleZen) is 13.37 %. The weighted average for players who play similarly to LittleZen is 13.29 %. So, he does 3-bet LittleZen less in the big blind than other players, but the difference is much much less substantial.

So, whereas LittleZen's play against knockstiff was so blatantly unusual that I was able to show immediately that it required further explanation, knockstiff's play vs. LittleZen

is a little unusual on its own, but only looks truly suspicious after a fairly complicated argument. It's not nearly as strong as my argument against LittleZen, especially when you consider that knockstiff's data is further corrupted by varying stack sizes (both knockstiff's varying stack size and the difference in the typical stack size of the opponents that I looked at).

However, LittleZen softplayed against knockstiff and presumably needed some incentive to do so. The data looks exactly like one would expect it to look if stoxtrader softplayed against LittleZen. It does not look like the play of an innocent reasonable player. So, the only reasonable conclusion is that the two players colluded.

However, I'm not completely comfortable with this. The argument is nuanced and complex. If stoxtrader would like me to change my mind, all he needs to do is provide me with all the hands he's played on Stars (e-mailed to me directly by Stars support, of course) and I will quickly determine conclusively whether or not he softplayed against LittleZen. I can of course assure him that I will only use this database to look into the collusion, and will delete it after I'm done. He's ignored these requests in the past, but hopefully this post will make him change his mind.

6. KINETICA

Kinetica is an FTP account that was accused of softplaying with stoxtrader's FTP multiaccount 40putts in March of '09, along with LittleZen and knockstiff on Stars. The fact that 40putts is owned by the same player as knockstiff, that players didn't know that at the time (to my knowledge), and still linked the two different collusion cases is pretty suspicious already. Now that I've shown that knockstiff and LittleZen did in fact collude, it's even more suspicious.

However, I don't need those facts. Kinetica's softplaying was quite blatant.

I analyzed data comes from CAP games on FTP. Only 41.9 % of Kinetica's hands where played at CAP tables, but CAP games are perfect for this investigation. In almost every

hand at a CAP table, every player has 30 BBs.⁶ This means that the problem of worrying about playing differently against different stack sizes or tending to have a different stack size when different players are at the table is completely gone.

30 BBs is a much more complicated stack size than the 20 BBs that LittleZen prefers, but it's still no where near as complicated as the 100 BB stacks that most of us play. Kinetica is a winning player in these games, so he understands this. Therefore the following data cannot be explained in any way except for collusion.

Figure 15 shows Kinetica's 3-bet % vs. his 49 most common opponent's open raises when it folded to him in a position other than the BB.


His 3-bet % against 40putts is 1 %. KK, AA, and AKs (not including AKo) is 1.2 % of hands. So he's presumably 3- betting even tighter than that range against 40putts on average. You don't need to know much about poker to realize that that's very strange.

The error bars are quite large because PTR's sample on Kinetica in just CAP games isn't very big. Like in the case of knockstiff, this results in the large variation in 3-bet %s that you see in the chart. It's pretty reasonable to assume that the two players who got 3-bet less than 3 % of the time are players who Kinetica actually 3-bets more frequently than that, but simply appear lower because of randomness. (The lowest value, 2.4 %, is also against a player with a 14.4 % VPIP in these games, compared to 40putts's 20.4 %.)

However, because Kinetica's play was so incredibly abnormal, I don't have to worry about all that. The odds that Kinetica's actual 3-bet % vs. 40putts was 2.4 % but PTR's database randomly showed 1% are quite low (about one in 3.1 million). If I exclude that data point because of the extremely low VPIP of that player and go to the next one, 2.8 %, the odds are extremely low (about one in 1.7 billion). If I use the weighted average 3-bet % (excluding 40putts) of 6.71 %, the odds are extremely low (about one in $3.0 * 10^43$, or. about one in 30 tredicillion).

⁶The betting is capped at 30 BBs and the min buy-in is 30 BBs. Most players just buy in for ≈ 100 BBs and almost never fall below 30 BBs, but it does happen occasionally.


FIGURE 15. Kinetica 3-bet % vs. Open Raises (CAP, All players, not BB, no callers)



However, this is 30 BB poker. On Stars the majority of the hands were played with at least one player under 25 BBs, so cold calling wasn't much of an option. (LittleZen only cold called knockstiff's open raise 2.7 % of the time when it folded to him in a position other than the BB in general. Only 1.1 % of the time with 25 BB stacks.) With 30 BBs, it's much more reasonable to cold call, so it's possible that Kinetica wanted to see flops with 40putts and therefore decided to call him quite lightly.

That's not the case. See Figure ??.

FIGURE 16. Kinetica Cold Call % vs. Open Raises (Cap, All Players, not BB, no callers)




Notes: The error bars on some of the points on the far right are very inaccurate because the approximation I'm using breaks down for such low samples.

Not surprisingly, Kinetica's cold call % varies wildly by player. A lot of this can be explained by how often he 3-bets (more 3-bets leave less opportunities for cold calls) and the player's PFR %s, etc. And a bit of it is of course just statistical variation. However, the important thing to note is that Kinetica does not cold call 40putts abnormally often. He cold calls his open raises 4 % of the time when it folds to him in a position other than the big blind (with a fairly small error of +/- 0.39 %). The weighted average of the rest

of the sample is 4.34 %. So, he actually calls him a bit less often than the average player (though this is not statistically significant).

Perhaps a better metric to look at is simply how many hands he plays against a player in general. (See Figure 17).

FIGURE 17. Kinetica Call or Raise vs. Opponent's Open (CAP, All players, not BB, no callers)



So, the picture's a little uglier here, but not too bad. Against the vast majority of opponents, Kinetica plays about 8-12 % of hands when they raise and it folds to him in a position other than the BB. Against a decent amount of them, he plays between 7 and 8 % or 12 and 14 %, and there are a couple outliers at the top and one at the bottom, which is to be expected. Against 40putts he plays only 5 %.

The next lowest value is 6.1 %. The odds that he actually plays 6.1 % of hands against 40putts opens (not in the BB with no callers as usual) but showed up as 5 % in PTR's

database due to random error are 1.2 %. So, it's not outside the realm of possibility. However, the 6.1 % player has a sample size of only 278 and the error value is 1.44 %. If we accept that he is likely an outlier because in 48 samples, there are bound to be a few players who are one or two standard deviation outliers, the player with the lowest value seems like a likely candidate. If we drop that value and move to the next (7 %), the odds are about one in 31,000. If we use the weighted average (excluding the data on 40putts) of 11.05 %, the odds are extremely low (about one in $39 * 10^{27}$, or 39 octillion).

Of course, this argument is less compelling than the case that just considered 3-bets. But this is because I took the extremely abnormal behavior (the low 3-bet %) and diluted it with some fairly normal behavior (his cold calling %). The fact that, after I do this, I'm still able to make a fairly strong argument that something abnormal is going on should completely put the rest the question of whether or not Kinetica is 3-betting 40putts so little simply because he wants to play flops with him (since he's playing less hands against him than against everyone else).

So, let's consider the other legitimate reasons that Kinetica might 3-bet 40putts so much less than everyone else. I'll go through this part quickly because I've already done it twice.

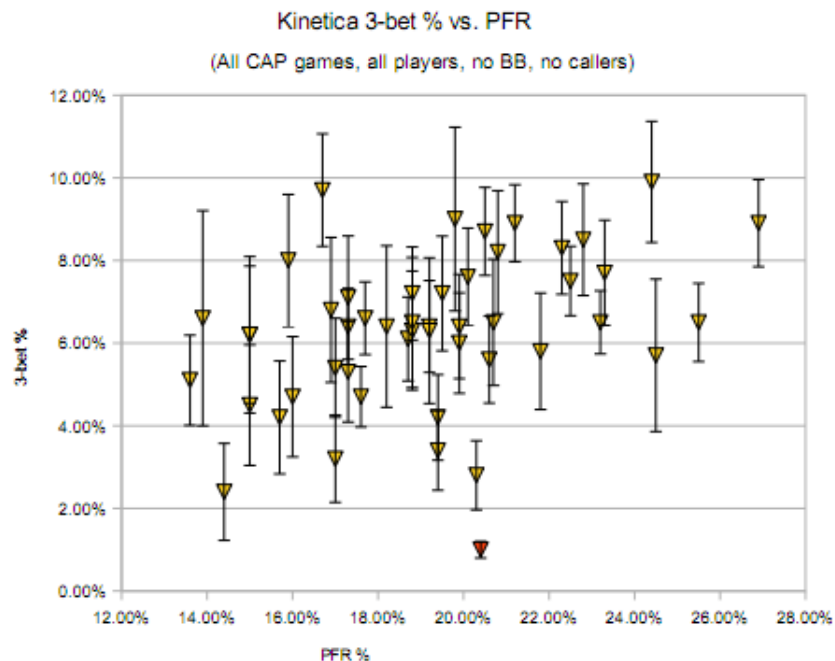
Figure 18 shows that 40putts actually raises fairly liberally.

Figure 18 shows that 40putts actually folds a lot to 3bets.

He opens fairly lightly and folds to tons of 3-bets, so he's a great candidate for 3-betting lightly, but Kinetica did not recognize this. How did Kinetica play against players who play similarly to 40putts? These data are shown in Figure 20.

Indeed, he played very differently. The lowest value there is 2.8 %. The odds that Kinetica actually 3-bet 40putts 2.8 % of the time but PTR got a value of 1 % due to random error are extremely low (about one in 1.7 billion). The weighted average of the sample is 6.88 %. The odds of that are even lower (about one in $1.4 * 10^{45}$ or about 1.4 quattuordecillion.).

FIGURE 18. Kinetica 3-bet % vs. PFR (All CAP Games, All players, not BB, no callers)



Maybe there's something I've missed. Maybe Kinetica doesn't like 3-betting 40putts for some legitimate reason that I didn't think of or maybe something irrational. If that were the case, he'd probably 3-bet him less in the big blind as well. If he's softplaying, he wouldn't have any reason to 3-bet less in the big blind (HU and closing the action).

The reader should examine Figure 21 to see that Kinetica clearly 3-bets 40putts about as often as he 3-bets everyone else from the big blind. He 3-bets him 10.5 % of the time in this spot. The weighted average against the other players in the sample is 11.43 %. So, I think I've made a very clear case that Kinetica went out of his way to avoid 3-betting 40putts. I've eliminated every remotely reasonable explanation I could think of. I've also shown that he only did this when a cheater would (not in the BB) and played normally

FIGURE 19. Kinetica 3-bet % vs. Fold to 3-bet (All CAP games, all players, no BB, no callers)

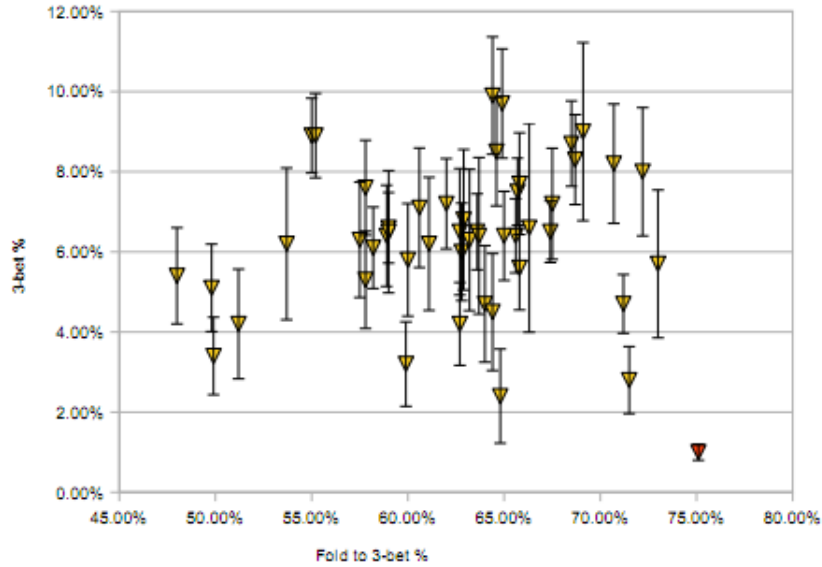


FIGURE 20. Kinetica 3-bet % vs. Opponents Similar to 40putts (CAP, PFRs within 1 % and F3s within 10 % of 40putts, no callers, not BB)





FIGURE 21. Kinetica 3-bet BB vs. 3-bet Not BB (CAP, All players, no callers)



Notes: These use different axes.

against 40putts in the most similar spot I could find where cheating wasn't possible (in the BB). So, Kinetica cheated.


7. 40PUTTS

As with knockstiff, the argument for 40putts is a bit less clear. However, it's not unreasonable to think that I don't need to make this case. Kinetica played in a way that benefited 40putts at his own expense. He's a winning poker player over a huge sample at mid/high-stakes, so he's not an idiot. It seems reasonable to assume that he was getting something in return for this.

I don't like this argument on its own. When I combine it with the data presented below, though, I think it's clear that 40putts softplayed Kinetica.

For 40putts, I also decided to look only at CAP games. The basic data are shown in Figure 22.

FIGURE 22. 40putts 3-bet % vs. His 49 Most Common Opponents (Not BB, No Callers)



A few things to notice here. As with all the other cases, his play against Kinetica is quite different than his normal play. He only 3-bet him 1.6 % of the time. For reference, KK+/AK is 2.1 %.

There is one extreme example of another player that I'm sure you've noticed. 40putts did not 3-bet this player once in the 122 relevant opportunities. The odds that an actual 4 % 3-bet resulted in this outlier are about .69 %. The odds that a 3 % 3-bet created it are about 2.4 %. Since this sample has 48 players (not counting Kinetica), it's not surprising to see some 2.4 % outliers. When you further consider that I've looked at two separate such samples on FTP and only found one piece of data that disagreed with the general conclusion, it's reasonable to assume that this might be an even less likely outlier. For example, a real value of 3.5 % has a 1.3 % (or one in 76) chance of resulting in this outlier.

Obviously, my argument to explain away the outlier this isn't the strongest. Were it not for the fact that Kinetica softplayed 40putts, I would probably stop here and say that things look suspicious but that I can't come to a clear conclusion. (In fact, I almost did that at one point when I looked at similar data on 40putts and knockstiff before looking at the data on LittleZen and Kinetica.) Given the situation with Kinetica AND the situation with LittleZen AND the situation with stox's stars account, knockstiff,

I think it's worth a further look.

Remember that I made a loose argument that suggested that the 0 % outlier might be in the 3 % to 3.5 % range. The odds that an actual 3-bet % of 3 % resulted in PTR's value of 1.6 % are about one in 19,000. The odds that an actual 3-bet % of 3.5 % resulted in PTR's value of 1.6 % are extremely low (about one in 4.5 million). Again, I want to stress that I got the 3 % and 3.5 % numbers from making an assumption based on a fairly loose argument. I believe that there's some justification for what I've done, but the argument is nuanced and a bit speculative, and I don't want anyone to just see that 4.5 million number and assume that that means conclusively that 40putts colluded. It's suggestive, but it's not concrete proof on its own.

However, let's see if breaking up the data will make things clearer. (See Figure 23).

The preflop raise helps a bit, but not much. You can see that in general stox seems to 3-bet people with similar PFRs to Kinetica much more often than Kinetica, so we know that Kinetica's PFR % isn't a satisfactory explanation for stox's odd play against Kinetica. But the 0 % player actually has a PFR quite close to that of 40putts. So he's still a problem.

The situation with the fold to 3-bet data is pretty much exactly the same (See Figure 24). Again, stox plays very differently against the vast majority of players with fold to 3-bets similar to Kinetica's. But again, the 0 % player can't be explained away by his preflop stats. He does fold to 3-bets a bit less often than Kinetica, but not significantly so.

FIGURE 23. 40putts 3-bet % vs. PFR (CAP, All Players, No Callers, No BB)






FIGURE 24. 40putts 3-bet % vs. Fold to 3-bet (CAP, All players, No callers, No BB)



40putts also does not make up for his lack of 3-bets against Kinetica by calling him more often (See Figure 25).

FIGURE 25. 40putts Cold Call % (CAP, All Players, not BB, No Callers)




Notes: The error bars in this graph are inaccurate for the players with low cold call percentages and low samples.

So this again doesn't explain 40putts's behavior against Kinetica. But again the player with the 0 % 3- bet has a similar result. The value for Kinetica is 4.5 %. The value for the outlier is 4.9 % (He's in yellow). He even 3-bets the outlier out of the BB about as often as Kinetica (See Figure 26).:

However, you'll also notice the familiar pattern of 40putts 3-betting Kinetica a reasonable amount of time from the big blind, which again is very suggestive of cheating.

At this point, I just asked PTR who this 0 % player is, and I really have nothing. He seems like a fairly normal mid/high-stakes grinder who played a bunch of hands with 40putts and happened not to get 3- bet much. I even checked to see if he's a multiaccount of Kinetica's, but he's played a bunch of hands with Kinetica, so that's not the case. The

FIGURE 26. 40putts 3-bet in BB vs. 3-bet not BB (CAP, No Callers, All Players)



fact that his play is pretty similar to Kinetica’s doesn’t tell us much either. Most people in these games play pretty similarly to Kinetica (except against 40putts, of course). I assume that some people will think that he’s another cheater. I highly doubt that this is the case because of how few hands he’s played with stox (about 2700). In particular, I don’t plan on releasing his name. Most likely, this number was just a random event that happened to make my life more difficult. However, I have no way to prove that that’s true.

In the case of knockstiff, I was able to show that the troublesome data points were explainable in two different ways. I showed that they could easily just be the result of statistical variation. I then also showed that these players played quite differently from LittleZen. When I looked at players who played

similarly to LittleZen, knockstiff showed a clear pattern that was only violated by LittleZen.

In 40putts’s case, things are a lot uglier. The only argument that I have is statistical variation. This is why I put 40putts last. If I saw this data out of context, I would’ve said that I thought he’d colluded, but I wasn’t totally sure because of this outlier. Add

in the context that Kinetica clearly softplays 40putts, and the case gets stronger, since Kinetica has no incentive to do this unless he gets something back from 40putts. Add in the additional context that stox owns 40putts and also owns the account knockstiff, which I've shown to be guilty of exactly this type of collusion, and I think it is clear that 40putts was also involved.

Again, this argument is a bit more complicated and nuanced than would be ideal. If stox would like to e-mail FTP support and tell them to send all hand histories played on his accounts to me, I'll gladly look them over and attempt to clear his name.