Electronic Trading: Winds of Change

Michael McElroy
Vice President, Independence Investment Associates

Evan Schulman
President, Lattice Trading

Introduction.

The subject of this talk is the application of electronic trading methods to enhance trading execution and performance measurement. It serves as a follow-up to the talk presented at the previous Northfield conference in Key West, Florida on transactions cost modeling.

First, electronic trading and its many manifestations will be defined, followed by a discussion of the systems available to implement these ideas. Next, a practitioner's view of implementation issues will be introduced, with particular reference to experiences at Independence Investment Associates. After laying this groundwork, issues in the evaluation of electronic trading data will be introduced, followed by quantitative methodology and the results of our short experience with real-time electronic trading data. Upon conclusion of this lead-in material, Evan Schulman will lead a discussion on the larger strategic directions facing electronic trading in the next five years, and we will conclude this talk with discussion of all the presented ideas.

Electronic Trading -- What is it?

Depending on who one asks, "electronic trading" is defined in many different ways. At the simplest level, a market structure such as the NASDAQ (National Association of Security Dealers Automated Quotation) system represents electronic trading, since multiple market makers are connected electronically into an network. Manual intervention is still part of this trading scheme, though -- a trader must determine the best bid and ask price via a terminal, and then place a call to the particular market maker (or use a service such as Instinet) in order to get an execution. While this market structure represents a higher level of automation than the NYSE specialist's "verbal" system, it still does not satisfy most current definitions of electronic trading.

The NYSE, through its DOT (Designated Order Turnaround) system, provides traders a conduit to direct orders directly to the specialist's post on the exchange floor, leapfrogging several levels of the information hierarchy associated with traditional trading. DOT represents a step in the quest toward electronic trading, for it allows traders to electronically send orders to and receive fills against these orders from the exchange floor; however, the one element DOT lacks is that it is only connected to one marketplace, or pool of liquidity. Other markets are also transacting contemporaneously, albeit "piggy-backing" off of the NYSE price-discovery mechanism, potential liquidity one is not seeing by using only DOT.

As mentioned above, there are also alternative structures. Crossing networks such as POSIT, Lattice, Instinet Crossing Network, and the Arizona Stock Exchange (AZX) also bring together buyers and sellers electronically with a structure similar to NASDAQ, although in a more automated fashion. These crossing networks, while not in the business of price discovery, provide alternative pools of liquidity for traders to take advantage of.

Finally, electronic trading (as I would like to define it for the rest of this talk) can be thought of as systems and strategies which poll and use these various aforementioned trading systems to provide best execution for the orders given to them. These systems might incorporate either short-term risk models, momentum strategies, or other algorithms to make the best decisions as to how and where to place orders. The design and measurement of these systems and strategies is going to be the subject of the rest of this talk.

These systems are available today to any buy or sell-side trading desk. DOT boxes, which connect the trading desk directly to the exchange specialist, can be obtained from almost any broker who is happy to provide the conduit" to the exchange floor, since the broker providing the box ends up making a commission on all trades done. The proliferation of DOT boxes is analogous to the competition among brands for supermarket shelf space -- there is limited real estate on trading desks, and having six boxes piled up (with six different communication protocols to boot!) is not the best use of this desk space, so DOT providers are now clamoring to provide extra features and lower cost trading.

Routing services (such as Davidge, Thomson Financial, Merrin, and others) attempt to consolidate these multiple boxes into one system. They provide a hub for getting orders to different brokers, while avoiding the problem of stacking a bunch of individual boxes on the desk. In addition, these routing services provide a common protocol to send orders to the router, which then provides the service of translating an order into the specific format(s) required by broker.

Crossing networks also provide alternative pools of liquidity. Vendors such as Instinet, AZX, POSIT, and Lattice provide crossing between buys and sell orders within the system, with the objective being to clear as many orders as possible subject to any client-specific constraints (cash constraints, pairs trading, etc.) One disadvantage of these crossing networks is that these are "discrete-time," meaning the cross happens either once or several times per day (in the case of POSIT) or after market hours (in the cases of Instinet and AZX), but is not continuously open throughout the day as the NYSE or NASDAQ markets are. This can create logistical problems for trading desks as they must prepare orders for crossing by removing them from the market, wait for the cross to happen, and then wait for a confirmation file to learn what the residuals are before re-commencing trading.

Lattice, has tried to provide a universal solution to this logistical headache. While providing multiple broker DOT access (as the routing systems have), it has also provided the capability for intra-day continuous time crossing in order to give the client maximum exposure to as much potential liquidity as possible. Just as with any of these systems (or markets, for that matter), the liquidity pool is dependent on how many participants comprise it. Developing a large and dedicated client base is a problem any of these alternative market structures face in competing with the NYSE or NASDAQ.

Electronic Trading -- How Does it Work?

Once these systems are installed on the desk, the work is not done. Orders must be generated, and trades against these orders must be processed, orders often residing on incompatible systems and platforms to the electronic trading.

DOT trading is fairly simplistic -- these systems can accept either "market" or "limit" orders only, which do not dynamically update as prices move. For example, if one has a strategy to "buy stock on the bid side, regardless of the price level" in a market which is currently 26 1/2 (bid) - 26 5/8 (offered), he/she would place a limit order to buy at a price of 26 1/2. If the price of the stock moves up so that the market then becomes 26 5/8 (bid) - 27 3/4 (offered), it is the burden of the user to cancel the 26 1/2 bid, wait for a confirmation from the specialist acknowledging the cancel, and then re-submit the new buy order at 26 5/8. Users of routing systems face the same problems, although they are provided with more broker choices and a single communication protocol with the different brokers. The burden of order refreshing still rests on the shoulders of the user.

Crossing networks (POSIT, Instinet, AZX) require only that a user submit his/her buy and sell lists, which then get agglomerated with other clients' lists into a master buy and sell list. When the discrete time arrives for the cross to take place (intra-day for POSIT, end-of-day for Instinet and AZX), the crossing systems take a snapshot of the market price (or close price) from the primary market -- most likely the NYSE -- and try to clear the lists of buys and sells. Taking the market price from the primary market is known as "piggy-backing" on the NYSE price discovery mechanism. Users of these crossing networks can also include account constraints along with their orders (cash constraints, such as must do an equal dollar amount of buys and sells), which clearly provide a sub-optimal solution for the overall optimization, but allow clients to reflect trading realities they face.

Lattice, as mentioned above, provides the best of both worlds. On the DOT side, Lattice provides for dynamic updating of orders, so a user can implement a strategy such as "I want to buy stock X on the bid side" and not have to worry about all the canceling and replacing that would be necessary with a simple DOT system. Additionally, the Lattice user can do continuous crossing with other users on the network as long as there are counterparties to trade with. Along with all this functionality comes the ability to use multiple brokers for clearing the trades, just as a routing system provides.

Electronic Trading Implementation at Independence.

In February 1993, we started implementing electronic trading technologies on our equity trading desk at Independence. There were several reasons why we did this. First, using these systems along with our traditional trading would help us to obtain best execution on our orders. Second, these systems would allow our traders to be liberated from simple orders (initially thought to be the only type of orders appropriate for these systems,) giving them more time to concentrate on the difficult, block orders. Lastly, these systems all have lower commission rates (ranging from less than 1 cent per share up to about 2.5 cents per share), so we would be able to reduce our commission costs.

The four systems we implemented were Lattice, AZX, Instinet and POSIT. This was not as easy as just plugging in a wire to the back of our computer network, since we had to integrate these systems with our existing mainframe-based order management system. Success at this required close contact with each vendor in order to understand software compatibility issues, post-trade billing and settlement procedures, and communications protocols between our networks and the individual vendors'. While much time was devoted to learning these systems and their compatibility requirements, we viewed it as time well spent -- it was a good way to get up the learning curve so that as other systems were developed and marketed to us, we would know exactly what it entailed to get them up and running.

Two other implementation issues were less concrete as we embarked on this project. The first was trading room "buy-in.". When we initially started the project we felt that giving the trading room the tools to trade electronically would be the best thing since sliced bread to them. Reality, in the first few months, turned out to be very different. For all those simple trades that used to be scattered around to the different brokerage relationships, suddenly there was no voice needed at the other end of the line, since these orders were being sent down to an exchange floor or crossing network automatically. Yet even worse than the potential for broken relationships was the perceived threat that electronic trading would eliminate the need for traders trading; rather, the traders felt they might be relegated to button pushing and file transfer duties. Clearly these were unfounded fears, yet I think this serves as a good lesson to keep in mind when implementing new technology in an organization.

The second unforeseen issue was the difficulty of implementing trading strategy in software. There were several issues to deal with here. First, we had to find out how we do it now, which was an interesting exercise to go through. Trading, which to this point had been somewhat of an arcane subject, was suddenly brought into the spotlight, and I am not sure the traders were ready to start enunciating how they traded. Even more threatening was the fact that rules which were used for traditional trading could be implemented in software, software which would provide economies of scale in terms of the number of names that could be analyzed and traded simultaneously. Lastly, our trading room became very worried about the performance of these algorithms versus the desk trading activity, which made it important that we be able to monitor and calculate performance numbers for trades done electronically.

There are several ideas to consider when implementing electronic trading strategies. First, the results one sees depend on the strategy employed. For instance, our strategy was to try and trade at the market price until the price moved against us, and then become more patient (limit) traders. Additionally, we wanted to make sure we finished the entire order if possible within the trading day. Using this strategy, the resulting trades will be a function of the benchmark price used to calculate market impact. For example, if we say the market price is the "wrong-side" price (ask price for buys, bid price for sells), and then we use as our benchmark price for trade performance measurement the "right-side" price (bid price for buys, ask price for sells), we will clearly always incur some market impact (unless the bid and ask prices are equal!) An even deeper question to answer is whether we can assume we will trade (either traditionally or electronically) at the right side price, or will we have to make a price concession in order to trade? The gist of this discussion is that the benchmark price one uses to calculate impact should relate to the mandate of the trading operation -- is the trading room considered a profit center (use the right-side price), or a cost center (wrong-side price)?

Another idea which becomes important in implementation of trading strategies is the definition of market impact -- should it be calculated on the ticket level, or on an aggregate book of business level? The use of electronic systems allows for a move away from traditional "trade blotter" and "average price" trading, so it must be decided how to measure impact -- locally, or globally.

There are a bunch of ideas to digest here, each with no correct answer. The benefits of successful implementation, though, are large. On an average priced stock of $50, reducing the market impact of paying a 1/8 point spread (which equal 12.5 cents) by half translates directly into 12.5 b.p. of portfolio performance enhancement. Taking actual trading results and trying to model the costs of trading (as the rest of this paper attempts to do) can lead to more realistic portfolio implementation in the trading operations.

Learning From Experience.

Electronic trading requires obtaining and warehousing a large amount of market-related data. This provides a rich playground for both quantitative analysts and data miners. Unfortunately, we cannot learn much from POSIT, Instinet, or AZX trades, since we do not have much information about them other than the orders we submitted and the trades we were successful on (with which hit ratios can be calculated.) These services (POSIT in particular) does not release overall attempted trading activity, nor do they release the spreads existing at the times the crosses happened. One must take it at face value that he/she is receiving the mid-point price of the market at the time of the cross. For the overnight crossing systems, the prices used are market closing prices, so there is not a whole lot of good information to be gleaned from this data. Trades done using the Lattice system, though, are rich with data primarily because taking complete control of the order in our own hands on the desk requires that we have a lot of initial condition and market data in order to place intelligently-priced and sized orders.

Slides #8, #9 and #10 show some data which was used in our study of explaining transactions cost. Slide #8 shows a scatter plot of over 14,000 orders submitted to the Lattice system over the past 14 months. On the horizontal axis are the orders, and on the vertical axis is the market impact (in percent.) As one can see, the average market impact (measured from the "wrong-side" price) is around zero, with tight clustering around the zero line. Market impact is a signed value, with negative market impact (buying below the ask price, or selling above the bid price) considered "good," and positive market impact considered "bad." Slide #9 shows the data behind the picture, with the variables explained on Slide #10. One variable, MKT_DEEP, needs explaining. It represents the "wrong-side" size (shares) minus the "right-side" size (shares), divided by the order size. For example, if we are buying 1000 shares of XYZ, and the sizes in the market are 5000 bid for, 20,000 offered, the market depth for this order would be (20000-5000)/1000 = 15. A positive value greater than 1 means that the market as shown is imbalanced against us, while a negative value means that we are going the same direction as the market. Values with absolute value less than 1 mean the markets are balanced among buyers and sellers. This measure can be used as an indication as to whether one should be a liquidity "provider" or a liquidity "taker" in the market.

With this data, we were curious as to whether we could learn any lessons applicable to trading. A regression of market impact (calculated from the previous night's closing price) versus the variables listed above was run, with the results shown on Slide #11. As can be seen, the only variable which turned out to be significant was the elapsed time of the trade, which was not really surprising. Even "worse" was the R^2 of 0.02. For those who saw the transactions cost modeling talk in Key West, these results are certainly on par with those presented there. There are two problems with this regression, though. First, using the previous night's closing price is not really a relevant benchmark, since one cannot transact there -- its only benefit is to measure the "slippage" of the trade relative to the valuation price. Second, traders really have a problem with measurement using this benchmark price, since they never really get a chance to transact there.

Recognizing these shortcomings, a second regression run was run using prices that existed at the time of the trade to calculate the market impact of the trade. This gives us a report card for the trading effort, rather than for the entire investment implementation system. Additionally, predictive ability of these results could be used to help in modeling our expected transactions costs during the portfolio optimization phase, which could potentially reduce slippage. The regression results (using market impact calculated from the midpoint price of the bid-offer spread at the time the order was placed) turned out to be much more explanatory. Spread as a percentage of stock price had an extremely significant coefficient, as did the order size as a percentage of 5-day trailing average volume and the elapsed time to trade. Market depth seemed to matter as well, though not as significantly as the other variables just mentioned. The regression yielded an R^2 of 0.27, much better than the 0.02 seen in the previous regression, but still not inspiring extreme confidence.

After seeing these results, and talking about them with Evan Schulman, he recommended that I look at the data broken out by the time of day that the order was placed. The table on Slide #13 shows the data set sliced into orders traded during different times of the day. The attempt was to see what variables offered the best explanatory power in each of the time buckets. A shaded box represented statistically significant regression coefficients at the 95% level. Some general conclusions can be drawn from this table. First, spread percentage always seems important, although this is a function of how one defines impact to begin with. Second, market depth matters at the beginning and the end of the day. This could mean that traders should provide liquidity at these times of day, since in times when we represent the minority direction of trading (as indicated by a market depth measure greater than one), this leads to a negative impact as indicated by the sign of the regression coefficient. Apparently, providing liquidity pays more in the morning than in the afternoon, but it matters at both times.

Observations.

First, there are several choices the buy-side has for electronic trading, ranging from the simple DOT connection to a more powerful and complex Lattice system. Using electronic trading systems, though, requires that the traders take control of the order onto the desk, rather than allowing it to "be worked" by a floor trader. In order to accomplish this requires **data** -- and lots of it. A nice natural fall-out of this is that the modeler has a rich and broad set of data with which to work.

Second, while the data sets used in this study were small, they seem to offer insight into the best style of trading at different times of day, and with different order characteristics. With more data, a functional form to explain transactions costs a priori might be able to be developed, although the data is very noisy. While it is too early to guess, perhaps cluster analysis would be a candidate method to use for this analysis.

An internal question one has to answer is whether altering trading strategy because of data analysis would still allow accomplishment of trading aims. For example, if a firm has a strong belief that residuals on the books are bad, but trading analysis says that in order to reduce impact one should become a more patient trader, how does this prescription conflict with internal policies? Should it matter, if market impact is reduced?

Overview

- Define Electronic Trading
- Systems Available and their Characteristics
- Implementation of Electronic Trading at Independence
- Issues in Evaluation of Electronic Trading Data
- Quantitative Methodology and Results
- Observations, Evan Schulman, and Discussion

Electronic Trading -- What is it?

- NASDAQ
- Designated Order Turnaround (DOT) trading
- Crossing Networks
- More Sophisticated, "hands-off" strategies incorporating above systems

Systems Available

- DOT Boxes

  Almost any broker will give you one

- Routing Systems

  Davidge, Thomson, Merrin, others

- Crossing Networks

  Instinet, Arizona Stock Exchange, POSIT, Lattice

- A Universal System

  Why not incorporate all these in one?

"Mechanics" of these Systems

- DOT

  Send specialist either market or limit order
  Not dynamic updating -- burden on user

- Routing Services

  Same as DOT, but more broker choices
  Vendor will provide single protocol

} Both send to exchanges

- Crossing Networks

  Users submit buy/sell lists, algorithm matches orders

  Separate "exchanges" piggy-backing on NYSE prices

- Lattice System

  Dynamic order updating
  Both DOT and crossing capabilities

Systems Implementation at Independence

- Which Systems?

  Lattice, Arizona, Instinet, POSIT

- When?

  February 1993

- How?

  Integrated systems with our existing mainframe order system
  Worked with vendors on understanding protocols

- Why?

  To obtain "best execution" in implementing investment philosophy

  Reduce commissions

Issues in System Implementation

- Integration of electronic trading systems with other systems

  Not many "turnkey" systems

  Good way to get up the learning curve

- Trading Room "buy-in"

  Suddenly, no voice at the other end of the line...

  Button-pushers versus paper pushers

- Implementation of Trading Strategy into Software

  How do we do it now?

  Oh, no, there goes my job!

  How does the computer do versus us?

Ideas to Ponder in Implementation of Trading Strategies

- Results obtained depend on strategy employed

  Our strategy: Try to trade at market price until price moves
  Try to finish entire order (no leaves)

- Results a function of benchmark price used to calculate market impact

  Can we assume buy on bid, sell on ask?
  Benchmark price should relate to mandate (profit/cost center)

- Impact calculated on "ticket" level, rather than aggregate trade level

  No more "average pricing"

- What are benefits of successful implementation?

  Reducing market impact by 6.25 cents = 12.5 b.p. in performance

  Market impact model = more realistic portfolio implementation

How Does that Computer Do? (Trades using Lattice)

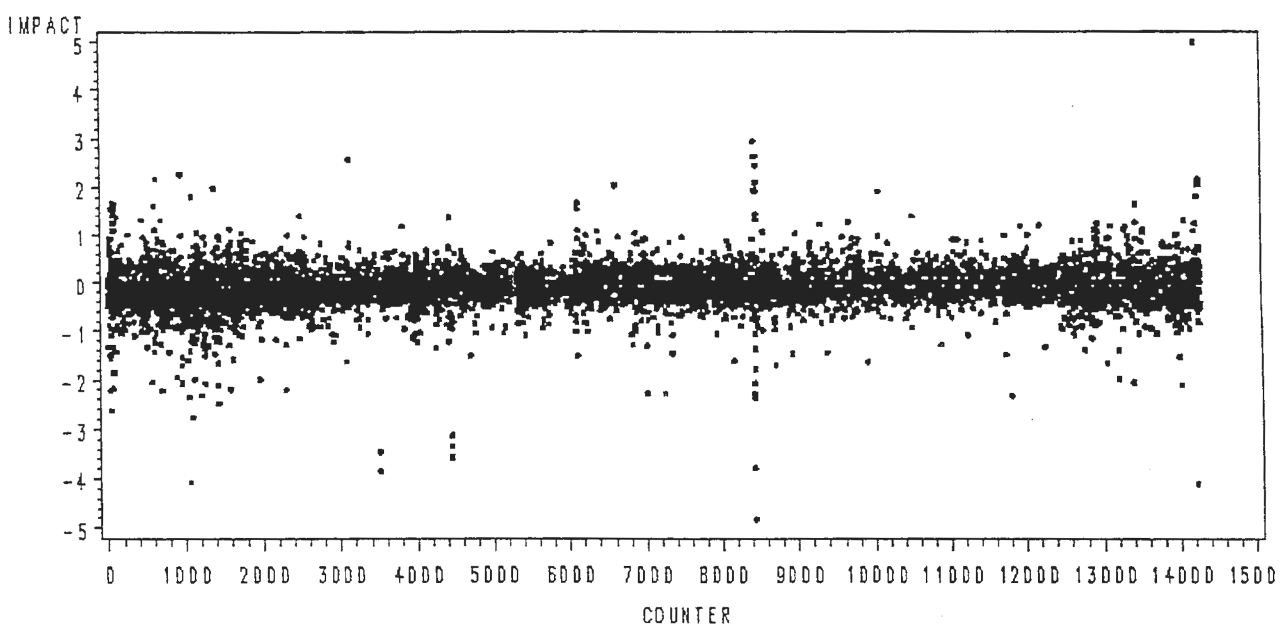

- **POSIT, Instinet, Arizona -- we have no control over!**

The Data Behind the Picture

	# obs	Mean	Std Dev	Min	Max
SHARES	14285	1616.19	2326.37	14	48600
AVGPRC	14285	55.792449	38.434178	4.25	681.875
IMPACT	14272	0.0015023	0.0032038	-0.046154	0.0456432
PCTMCAP	14192	1.184E-05	2.344E-05	3.73E-08	0.0005638
SPDPCT	14284	0.0043485	0.0027519	-0.008299	0.0454545
PCTAVOL	14223	0.0049093	0.0160178	1.209E-05	1
MKT_DEEP	11612	2.0180208	45.89546	-635	989
ELAPSED	2352	962.19898	856.0508	247	12883
TURNOVER	14192	0.0031495	0.0023098	1.333E-05	0.0328798

Details, Details... (modeling issues)

- Impact = difference between average trade price and benchmark price

  Two different benchmark prices used: Previous close & midpoint

  Positive is bad, negative is good

Variables used to explain market impact

Percentage of market cap order represents (PCTMCAP)
Spread as percentage of price (SPDPCT)
Shares as percentage of trailing 5-day avg volume (PCTAVOL)
Market "Depth" (MKT_DEEP)
Percentage of volume over time of trading (PCTVOL)
Elapsed time of trade (ELAPSED)
Time of Day

Market Impact Regression (Previous Close)

Explain calculated impact as function of aforementioned variables

Not really relevant benchmark -- can't transact there!

Traders grimace at the thought of being measured by this price

Tough to explain market impact -- R^2 of only 0.02 (just like last year!)

Parameter Estimates

Variable	DF	Parameter Estimate	Standard Error	T for H0: Parameter=0	Prob > |T|
PCTMCAP	1	-12.647842	16.70281600	-0.757	0.4490
SPDPCT	1	0.035071	0.08173180	0.429	0.6679
PCTVOL	1	0.000431	0.00038277	1.125	0.2608
PCTAVOL	1	0.025514	0.04790852	0.533	0.5944
MKT_DEEP	1	-0.000005326	0.00000570	-0.934	0.3503
ELAPSED	1	0.000001376	0.00000030	4.548	0.0001

Market Impact Regression (Trade-time Prices)

- Explain market impact calculated from prices existing at time of trade
  - *A report card for the trader, versus one for the organization*
  - *If "explainable," use results in real-time optimization routine*
- Much better results explaining time-of-trade impact -- R^2 of 0.27

Parameter Estimates

Variable	DF	Parameter Estimate	Standard Error	T for H0: Parameter=0	Prob > |T|
PCTMCAP	1	-0.664317	3.75952278	-0.177	0.8598
SPDPCT	1	0.188818	0.01839645	10.264	0.0001
PCTVOL	1	-0.000023458	0.00008615	-0.272	0.7854
PCTAVOL	1	0.051585	0.01078340	4.784	0.0001
MKT_DEEP	1	-0.000002773	0.00000128	-2.161	0.0308
ELAPSED	1	0.000000684	0.00000007	10.045	0.0001

The Next Level of Detail -- Time of Day

- Slice data set into orders traded during different times of day

  Seven slices: 9-10, 10-11, 11-12, 12-1, 1-2, 2-3, 3-4
- See what variables offer best explanatory ability during different times

	9-10	10-11	11-12	12-1	1-2	2-3	3-4
PCTMCAP	5.89	-7.99	1.73	-11.34	-16	-2.78	132.1
SPDPCT	0.243	0.283	0.071	0.281	0.187	0.246	-0.04
PCTVOL	-3E-04	-8E-05	3E-04	2E-05	0.001	3E-04	5E-04
PCTAVOL	0.039	0.078	0.063	0.077	0.043	0.034	-0.35
MKT_DEEP	-1.57	-0.15	0.13	-0.13	-0.38	-0.57	-0.59
ELAPSED	9.89	-0.35	9.41	2.67	2.71	6.1	20.2
R-squared	0.326	0.205	0.266	0.342	0.229	0.38	0.347
# Trades	132	422	491	517	280	317	119

Observations

- Electronic trading forces the discipline of gathering all sorts of data

  By-product: more detailed data for modeling
- Though data sets small, offer insight into best style of trading

  With more data, could potentially generate functional form

  Very noisy data -- perhaps cluster analysis of results?
- If trading strategy altered, can we still achieve our aims?

  Is it necessary, if we reduce impact?
- Other ideas, other data to get...