



Lecture 8: Strategies When there is no Strategy: Minority Game



DEMOLE A. C. - Zig, Zag, Zoug! Recueil de jeux pour la jeunesse. Delachaux & Niestlee, 1943 and 1946.


Zig-zag-zoug is a Suisse children's game, where you either put your foot in or out. Being in the minority wins. See the Zig Zag Zoug band on <http://www.youtube.com/watch?v=DwZC4yZoH4U> OR watch the Swiss army play the game <https://www.youtube.com/watch?v=28laQF0-Gyw>

The minority game has $2N+1$ players who choose one out of two alternatives. Being in **minority** wins. The MG is based on Brian Arthur's El Farol Bar problem. *El Farol has one of the best bars on earth!* NYT Dec, 2005



You want to go to the bar if $< 60\%$ of the population go; but you don't want to go if more than 60% go. Your decision depends on your expectations of how many other people will go ... but the decisions of the others depends on their expectations, so your decision depends on your expectation of their expectation."The El Farol problem is a fierce assault on the conventions of standard economics."

Brian Arthur simulated the attendance of 100 possible bar goers over 100 weeks. No goer would attend if he or she predicted that the crowd would exceed 60 (the 'threshold'). Over time, the number of attendees, which appears to be a random (or chaotic) sequence, fluctuates around the threshold.



If everyone used the same rule – **a map between past outcomes and their decision** -- either everyone would go or not go – herding phenomenon would produce worst outcome. So to explain pattern, assume heterogeneity of **decision rules** – say 60% look at the past outcome (could be too crowded or too empty) and their rule says go while 40% have a decision rule that says don't go. Then with some random variation you would get around 60% going. Alternatively, every decision rule could be probabilistic that if too many went last period, go next time with 70% chance – or 30% chance, etc. and if too few went last period, go with 50% or 90% chance,

Consider a “zero intelligence” rule where people pick an integer from 1 to 100 without replacement (if someone picks 50, it is no longer available for next person to pick) for their estimate, then always get 60 people predicting 60 or less and will 60. If predictions are with replacement on average 60 people will predict < 60 and go; and each period 60% will go. The variance of attendance will be variance of binomial — $100(.6)(.4) = 24$, so often will get higher than 60 – inefficient, with everyone unhappy. Or lower than 60% , with some happy but inefficient as more could be happy going.

But is $< 60\%$ going the correct utility? Do you want to go to bar and be only person there? More reasonable utility is that everyone want lots of people there but that the bar be not crowded. The landscape

for the El Farol problem would then be single-peaked around 60 with utility falling off as number going was deviated from 60, with 59 being pretty good and 0 or 100 being terrible.

Can some intelligence reduce variance?

1) There is no deductively rational solution since agents do not know expectational model of other agents. If someone knows everyone else's rules, they could see what people would do based on past history, determine how many would go and go if number was below 60 or stay home if it was above 60.

2) Expectations have divergent property: if all believe most will go, nobody will go, invalidating that belief.



3) Intrinsically dynamic and historic because people will always reassess their chances of success.

Minority game generalizes El Farol with different reward structures/utility function. Game has $2N+1$ players – odd number to eliminate 50/50 outcome. It is non-zero sum.

A) Each player in minority gets 1 pt, then social optimum is for N people in minority and $N+1$ in majority: **7 players**, 3 winners. Output is 3 if 3 on minority; 2 if 2 are on minority, 1 if just 1 is minority. People are rewarded for not herding.

B) Tournament where each player gets $G_i = -A_i \sum A_j$, where $A_{\text{sub}} = 1$ or **-1**, and **A_j is for all players except i** . When a player meets someone like him/herself loses -1 (- x -) or -1 (1 x 1) while gains 1 from meeting player unlike self (-)(+)(-). Gain for all is $\sum G = \sum (-A_i \sum A_j)$.

If all players pick 1 (or -1), each loses 6, for a social outcome of -42

If one player picks 1 and others pick -1 the “winner” gets 6 but 6 losers get -5 so social outcome is -24.

If 2 players pick 1, each gains 4; while the 5 players in majority lose 2 each for a social outcome is -6.

If 3 players pick 1 each gains 2 while 4 losers gain 0 (-1 when they meet another -1 player; 1 when they meet a + player) for a social outcome of 6

Thus, we have -42, -24, -6, 6, 6, -6, -24, -42 as # who pick 1 goes from 0 to 6. **Higher value closer to balance.**

The MG reflects ANY choice situation where players interact through a single price that depends on supply-demand market balance:

Career choice — you want to run for Prez but not if the market is flooded with candidates.

Stock market — you want to sell when buyers prevail; buy when sellers prevail.

Queueing — which lane to pick; time to drive; airport lane for screening; Party/Room with other gender

Try <http://ccl.northwestern.edu/netlogo/models/MinorityGame>, where you can choose your strategy.

Physicists use mean field theory of statistical mechanics where the aggregate (not near neighbors) properties/determine decision to study the minority game: "The MG is a complex dynamical disordered system which can be understood with techniques from statistical physics." using spin-glass techniques (<http://arxiv.org/pdf/cond-mat/0205262.pdf>) but most exciting outcomes come from simulations and experiments with people.

Key assumptions

1 — Common knowledge of outcomes — all agents know the minority winners of last M periods.

2 — Learn from own experience; not know anything about the others; no collusion

3 — Agents base decisions only on information they have – how **strategies in their portfolio** did.

4 — **Heterogeneous** strategies since individuals convert info into expectations differently → different rules.

The MG is about interaction between agents and information -- intrinsically dynamic — adaption and learning with no settling down since each period, you get new information that leads you to another choice. Expectations

and behavior are at odds: if everyone thinks 1 wins, they play 0, but if they think that everyone thinks like them, they play 1 but if ... Worst outcome is if everyone uses same strategy. Optimal is near to half/half: bookie setting odds for bets.

There are Nash equilibrium in MG. If people are doing a mixed strategy, best is to choose 1/2 go/not go. Then over long run everyone will have same chance of being in minority – about 1/2th the time. Note that with N people on Sell and N+1 on Buy, person from Buy who switches while no one else changes loses; (if person from N switches and no one else changes, ditto). So better to stick with choice. But if more than one shifts, switchers can win: 2 people in minority SELL think lots of people in majority BUY will shift and lots do, the two shifters from BUY to SELL will win while the “lots” who go other way lose again. What is lots?

MG models decisions “as if agents look for patterns” in history and “predict” next move from how rules did in the past — **inductive rules**. But all that matters is a rule/mapping that takes past history to a decision – often a **look-up table** which tells you what to do as a function of last M outcomes. If putting foot in (1) won last 3 times and one of your two decision rules is 111—> 1 and another is 111-->0, you will go 1 if the strategy to which 111-->1 is a member did better in the past. The story is in the rules.

What are these decision rules? Agents are allocated S (=2 in all models) strategies from the population of all strategies. They compute how their strategies did in the past and use the highest payoff strategy to choose in the next period. Number of outcomes M has 2^M possible histories. The row below shows one response to all 8 histories when M =3. STRATEGY S1

history	0 0 0	0 0 1	0 1 0	1 0 0	0 1 1	1 0 1	1 1 0	1 1 1
action	1	0	0	1	0	1	0	0

The decision-maker chooses 1 or 0 in response to each history, there are 2⁸ = 256 possible strategies. What would other strategies be? Choose 1 for every case. STRATEGY S2

history	0 0 0	0 0 1	0 1 0	1 0 0	0 1 1	1 0 1	1 1 0	1 1 1
action	1	1	1	1	1	1	1	1

Choose as in first strategy but 0 in 000 case. STRATEGY S3

history	0 0 0	0 0 1	0 1 0	1 0 0	0 1 1	1 0 1	1 1 0	1 1 1
action	0	0	0	1	0	1	0	0

Observe the actual history ... say it was 011. Strategy S1 says choose 0. Strategy S2 says choose 1. Strategy S3 says 0. If S1 outscored S2 and S3 in the past, use S1 and choose 0.

The model has N players and M periods. The other potential parameter is S, the # of strategies in players' set of possibles. But NS enters analysis symmetrically, so no loss if S=2 and N varies. **Longer histories produce more strategies:** For M=2, 4 histories, and 2⁴ = 16 strategies. For M=3, there are 8 histories and 256 strategies (2⁸). For M=5, 32 histories and > 4 billion strategies (2³² = 4,294,967,296). Half of strategies are 'polar opposites" that transform every 0 to 1 and 1 to 0.

Optimal Outcome has ~ 50% choosing 0 and 1. The efficiency of a solution is its variance relative to N. If code decisions as -1 or 1 so that mean is 0. With this metric

	RANDOM			IDEAL COORDINATION			HERD: DO SAME		
#choosing 1	prob	value	(value-mean) ²	prob	value	(value-mean) ²	prob	value	(value-mean) ²
All 3	1/8	3	9	0			0.5	3	9
2	3/8	1	1	0.5	1	1	0		
1	3/8	-1	1	0.5	-1	1	0		
None	1/8	-3	9	0			0.5	-3	9
Σ (p (value-mean) ²	(9+3+3+9 =20) /8 = 3			(1+1)/2 = 1			(9+9)/2 = 9		

Social value measured as variance relative to N. If **RANDOMLY** choose strategies get a symmetric

distribution with an average mean of 0 and standard deviation of 1 – normal as N grows. Very unlikely to get all feet in/out with large N . $\sigma^2 = N$ so $\sigma^2/N = 1$.


If **IDEAL** Coordination $\sigma^2 = 1$ so $\sigma^2/N = 1/N \rightarrow 0$ as N gets large

If **HERDING** get bifurcated distribution. If one strategy is a leader and others follow as in straight-line CA model where people follow criminal or good guys MG ends up on the ends. $\sigma^2 = N^2$ so $\sigma^2/N = N$ which goes to infinity as N gets large.


Key MG result: Basing decisions on induction from past performance of strategies yields outcomes close to ideal coordination, with narrower dispersion than random.

Why does inductive selection of strategies produce a balanced market?

Solution like this



Not like this



1. Strategies and selection rule for choosing one strategy over other get closer to coordination than to random because individuals keep changing and adapting as history develops.

2. Simulation and analysis shows that solutions depend on **# of histories relative to # of players by an explicit parameter, $\alpha = 2^M/N$** , which measures how well players cover the "history/possible strategy space".

At **low α** , many players N relative to histories with result that many use same strategy and operate as crowd $\rightarrow$ "herd effect" $\rightarrow$ inefficient outcome worse than random. Further, likely that some agent will profit from any predictable pattern and their activity will wipe out that pattern, producing symmetric phase.

With high α lots of histories per person so choices are close to random with result that average outcome is low but better than herd outcome.

3. If you have "exogenous information" false histories and announce numbers get the same qualitative results (but less fluctuations). The phenomena are the result of **collective processes by which agents respond differently to the same information (whether correct or not)**. It doesn't matter if A won last 3 periods, just tell people A won and get same outcomes but more variation in area where people switch a lot.

4. There is an optimal α . If you start with very low α and increase it you reduce the herd behavior and thus get better solution – namely lower variance around mean. In other direction start with very high α so you are close to random and you lower α , you move more to coordinated outcome and reduce the variance.





Figure 3: Volatility as a function of the control parameter $\alpha = 2^m/N$ for $s = 2$ and different number of agents $N = 101, 201, 301, 501, 701$ ($\square, \diamond, \triangle, \nabla, \triangleleft$, respectively). Inset: Agent's mean success rate as function of α .

5. But efficiency is associated with high inequality! Calculate Gini coefficient for distribution of incomes under different values of $\alpha \rightarrow$ **inequality rises with efficiency. Why?**

When α is small get low inequality because many people use same strategy so similar incomes.

When α is large also get low inequality because most strategies are uncorrelated so it is random choices.

Near the 0.34 point, inequality is highest because best strategy of different people nearly always loses or wins → wide spread of incomes. That inequality is great near the point of maximum implicit cooperation among players disturbs some analysts but pleases other folk.



But “the rule of the original form of the minority game has an essential defect. The winners receive the same reward no matter how small the number of the minority agents is. This is unnatural, and this is the very reason why the conventional form of the minority game fails to reproduce the realistic wealth distribution among agents, namely, Gini's coefficient being much too small compared to the observed value. To remedy this, winners with larger risk ought to be rewarded better...this modification saves the model and reproduce the realistic wealth distribution known from the statistics.” (Minority Game as a Model for the Artificial Financial Markets Tanaka-Yamawaki, M.; Tokuoka, S. Evolutionary Computation, 2006. CEC. IEEE Congress on Pp : 2157 – 2162). So having more sophisticated and sensible-seeming reward system ---> higher inequality.

If agents choose good but not top strategy, get better results — "thermal minority game". Similarly, get better results by adding noise. Adding randomness improves situation when you have implicit herding.

6. Symmetry, asymmetry and cycles. Consider $N=101$ and $M=1$, which has 2 histories and 4 strategies: $S_1, 0 \rightarrow 1 \rightarrow 1$; $S_2, 0 \rightarrow 1 \rightarrow 0$; $S_3, 0 \rightarrow 0, 1 \rightarrow 1$; S_4 is $0 \rightarrow 0 \rightarrow 0$.

Each person has 2 strategies with 6 possible pairs so when α is small – ie have many people relative to strategies, many will have same strategy and move same way whatever the outcome. Each strategy alternatively wins and loses when same history appears --> giving approximately the same win probability for any history. **Symmetric phase.** But can also get cyclic cobweb style functions. With $M=1$ people respond to last period. Long memories rational expectations much less pronounced fluctuations.

If α is large there are more histories than players so strategies are largely uncorrelated. It is as if each player is making random choices in the game. This resembles random. This is **asymmetric phase**.




Figure 3.5: Histogram of probabilities $P(1|\mu)$ of winning action to be "1" given the history string μ . Plots are decimal representations of binary strings of information. $N = 101$ agents and $S = 2$ strategies with a linear payoff scheme. a) symmetric phase with $M = 5$ ($\alpha = 0.316$) b) asymmetric phase with $M = 6$ ($\alpha = 0.634$) [6]

Conclusion

1. INDUCTIVE REASONING from complicated space gives solution with adaptive agents “trying to predict” next winning choice, which is determined only by their own choices.
2. Model is "technical trading strategies" not production/trade because not all agents can win at the same time. A betting market, where "efficiency" is with half of the money on one side and half on the other side.
3. Some models add people playing other strategies – "producers" who respond to real news while speculators play minority game. Must have speculators to generate anything like real market outcomes.
4. Brian Arthur: “El Farol emphasized (for me) the difficulties of formulating economic behavior in ill-defined problems. The Minority Game emphasizes something different: the efficiency of the solution. This is as should be. The investigation reveals explicitly how strategies co-adapt and how efficiency is related to information. This opens an important door to understanding financial markets.” Markets can equilibrate with little intelligence and can screw up because smart people manipulate them.
5. Mark Buchanan: “Qualitatively, the minority game is a little like a market. It's a situation in which each person tries to profit by choosing the right strategy, but where what is "right" is determined by the collective actions of everyone together. There is no "right" independent of what others do. The minority game, like any market, is a game involving an enormous and complex ecology of interacting strategies. ... even without any external information hitting the market ... the price fluctuates up and down unpredictably merely through the perpetual evolution of the agents' trading strategies. One day, apparently caused by nothing, the price may suddenly drop by 10% -- all because many agents chose to sell just as a part of trying to predict what was about to happen next.

What about using reinforcement learning to play minority game?

Q learning in the minority game M. AndreCut and M. K. Ali Phys. Rev. E 64, 067103 – 26 November 2001 We present a numerical investigation of the minority game model, where the dynamics of the agents is described by the Q-learning algorithm ... that... show that the Q-learning dynamics is suppressing the “crowd effect,” which is characteristic of the minority game with standard inductive dynamics, and it converges to a stationary state that is close to the optimal Nash equilibrium of the game.

Reinforcement learning meets minority game: Toward optimal resource allocation Si-Ping Zhang, Jia-Qi Dong, Li Liu, Zi-Gang Huang, Liang Huang, and Ying-Cheng Lai Phys. Rev. E 99, 032302 – 6 March 2019 When agents are empowered with the popular Q-learning algorithm in AI ,, (in an) unknown game environment gradually and attempt to deliver the optimal actions to maximize the payoff, herding can be eliminated.... the system evolves persistently and relentlessly toward the optimal state in which all resources are used efficiently. However... large fluctuations occur intermittently in time. The statistical distribution of the time between two successive fluctuating events depends on whether the number of time steps in between is odd or even. . Since AI is becoming increasingly widespread, we expect our RL empowered minority game system to have applications.




Figure 4. Theoretical prediction of the variance σ^2 in comparison with the simulation results. The system has size $N = 1001$ and power-law degree distribution $P(k)$ with scaling exponent $\gamma = 3$. The theoretical prediction does not depend on the value of the average degree. In direct simulations, the values of the average degree are $\langle k \rangle = 6, 10, 14$, and 40 . The simulation results denoted by symbols are the same as those plotted in Fig. 1, with the pinning pattern indicator to be $\eta_+ = 0.8$.

Controlling herding in minority game systems Ji-Qiang Zhang¹, Zi-Gang Huang^{1,2}, Zhi-Xi Wu¹, Riqi Su² & Ying-Cheng Lai Spontaneous evolution of the resource allocation dynamics, however, often leads to a harmful herding behavior accompanied by strong fluctuations in which a large majority of agents crowd temporarily for a few resources, leaving many others unused. We ... find the universal existence of an optimal pinning fraction – the fraction of agents chosen to hold a fixed state, and the “pinning pattern,” the configuration of plus or minus state assigned to the pinned agents. to minimize the variance of the system, regardless of the pinning patterns and the network topology. Papers China + Arizona State