Two main problems arise from this view.

The amount of resources devoted to winning the lottery for political favours are too low for the size of the favours. If this is truly how winners of political favours are determined, then politicians are being highly irrational by not extracting larger tributes from the lottery participants when they are able to. It would be like repeatedly running a lottery with a \$100 prize and selling just \$5 worth of tickets. It makes no sense, and still leaves \$95 of the political gift unexplained, suggesting there is still a missing element at play in political favoritism.

If political donations are a way to purchase lottery tickets for political favours, then the way donors choose to split their donations amongst political parties should merely follow political polling as donors seek to ensure they are entering a lottery with a genuine future payoff. Yet this is not the case at all. The image below shows the frequency distribution of donation shares between major parties amongst political donors in Australia (and its States), Canada, UK and Germany.

The blue line is the frequency of donations at each split between parties. Looking at the top left panel we see that around a third of political donors give 100% of their donation to the Labor party in Australia, while around half donate exclusively to the Liberal National Party (LNP). However, there is a small group of donors who donate equally (at 0.5) to each party. We see a clustering of donors who split their donation between major parties 50:50 in Australia and Canada, and at 60:40 in Germany, along with the lack of clustering in the UK. The orange line weights this frequency by the monetary size of the donation by each donor. Donors who split their donation 50:50 donate more money than those who are loyal to one party of the other, as seen by the higher orange line around the 50:50 mark. Such patterns simply cannot be predicted by rent-seeking theory.

From an evolutionary perspective however, these problems fade away. We no longer have a world of politicians and regulators seeking to extort the optimal price in a lottery for their political favours. Nor do we have a world of irrational regulators cooperating with former employers when they should instead defect.

What we have is world of tribe formation through the reciprocal exchange of favors. Political donations can be seen merely an entry fee to these tribes; an identifying marker, or to borrow a gang analogy, a hard-to-fake "facial tattoo" signalling future loyalty. In keeping with this reasoning, politically stronger groups may have higher entry fees, which can explain the dominance of a 60:40 split in Germany. Rather than an unexplainable curiosity, the clustering of political donations split 50:50 simply indicates these donors are paying entry fees to roughly equally powerful political tribes on both sides of politics. The majority of donors remain members of one political tribe only, donating exclusively to them regardless of political polling.

Looking at the revolving door through the evolutionary lens of group formation we get a real sense of the underlying conundrum expressed by Clinton. Do regulators revolving into positions of power already have implicit loyalties to their tribe, providing back-scratching favours to them when possible? And can

these loyalties be overcome in order to act in the interests of the public at large, even when it comes at a cost to their tribe? Should we also be looking for other signals, or identifiers of loyalty to one tribe or another from key personnel in regulatory agencies? These are important questions for evolutionary social scientists.

Studying back-scratching in the wild faces a major problem. Political favors are almost impossible to objectively observe. Not only is there a powerful incentive to conceal favors, determining the 'no favor' counterfactual is almost impossible. Was the government contract given to the most efficient firm? Or was it a favor to the winner, which we know for sure because we have information about an alternative bidder who could have delivered a better outcome for the price? More often than not we just don't know.

What makes back-scratching costly to the group as a whole is that each of the potential recipients is given a randomly shuffled 'productivity number' each round, either 1, 2 or 3, which determines a payoff for everyone in the group (see the image below). Each player gets an amount equal to the receivers productivity number in a round. Give the prize to the player who is 1, the group gets \$1 each. Give the prize to the player with 3, the group gets \$3 each. Think of this productivity number as reflecting the efficiency firms competing for a government contract.

To sustain a back-scratching alliance means not choosing the most productive player for the prize in two-thirds of the rounds. It earns the alliance \$725, the outside group \$100, and comes with an efficiency loss of \$100 to the group of four as a whole due to the repeated non-productive choices, which is shown below as the emergence of a loyal sub-group.

It turns out that in 84% of groups formed costly tribe-like alliances at some point, even though the participants were young honest university students who were anonymous. In real money terms (rather than experimental currency) alliance pairs 'stole' \$30 from the outsiders to increase their alliance earnings by \$20.

I also formed groups where I primed subjects before the game. They read a brief recount of their imaginary 'in-game' relationships with a connected partner, effectively being told that one of the symbols was their friend that helped them in the past. As you would have guessed, this increased the number of rounds of favours substantially; from 28% to 62% of all rounds.

Surprisingly alliance players were happy about their actions. They felt like they were being cooperative and social, and didn't feel guilty about the costs they imposed on others. They also rationalised their behaviour, saying that forming an alliance is a justifiable strategy, while also concealing it by lying when explicitly asked in a later survey if they had formed an alliance. These insights are important when considering how to combat back-scratching in the political world; people will attempt to delude themselves about the morality of their actions.

I tested two institutional changes in the experiment after the initial 25 rounds that were designed to combat back-scratching. The first was a randomization mechanism that picked the next round allocator of the gift with a degree of randomness, and was designed to mimic policies of staff rotation, which are

standard in areas such as auditing and customs. Think of the this as random choice of new staff rather than a choice of staff with industry experience who are revolving into the organization. Second was a policy mimicking bureaucratic procedures to limit the size of favors able to be allocated with discretion, which decreased the gift each round from \$25 to \$3, making the payoff from an alliance lower than not being in one. This low-rent treatment I describe as a McDonalds-ification of decisions; take away discretion for decision makers, such as politicians and regulators, and have them follow a set procedure for major decisions. In a McDonalds bureaucracy I can't give away free burgers to my mates, but at my uncle's burger shop where I have supreme discretion, I can give away burgers with near impunity.

Somewhat surprisingly, the staff rotation policy didn't curtail back-scratching on average because it allowed two competing alliances to take temporary control and make favours. It was a lot like choosing a referee from two competing football teams at random; the referee stayed loyal to their team and their favours still made the game inefficient. However for those groups who had not yet formed strong alliances, it did deter them from forming.

The opposite was the case with the low-rent policy. Those groups who didn't yet have alliances emerge from favouritism saw them start, possibly because the cost of doing so to others had decreased as well. Those groups with strong alliances did see reductions in favouritism and an improvement in group outcomes. This could be a reason organisations seeking internal fairness do routinely implement such extensive policies and procedures, because it does stop friends doing favours within the organisation.

Other notable observations are that social norms are very strong. In the experiment, like we see in real life, groups where some people are observed 'doing the right thing for the group' quickly lead to this becoming the norm. Whereas where favouritism is observed in early rounds, the group descends into counterproductive back-scratching.

Another observation is that it pays to be loyal, but not too loyal. If your alliance partner fails to come through with favours it pays to look for someone else whose back needs scratching.

None of this is really rocket science. If this sounds obvious to you then you've probably learnt a lot about human behaviour through life experience. It also means the experiment is capturing some critical elements of real phenomena.

Anti-corruption policy translated from the lab

Policies to tackle the revolving door in Washington is the real value here. Some that immediately come to mind based on these experimental results are:

Independence requirements. While many regulatory agencies feel the need for staff to have industry experience, this experience comes with a set of alliances to others. However, if staff come from other countries (or states) they may still have industry experience, but no tribe-like relationships to the people in the companies being regulated.

Cooling-off periods. The flip-side of this is to have quite long cooling-off periods for government and agency staff. The intention is that during this time the value of relationships depreciates as new people and policies emerge in different organizations, changing the structure of any insider group and reducing the ability to repay favours with future favors.

Internal rule systems. Reducing the amount of discretion that any particular employee has in a regulatory agency, even at high levels, will ensure that any favoritism requires the cooperation of a larger group. Voting systems that require high levels of approval, multiple levels of oversight, and transparency of information flow and decisions, are simple and possibly effective interventions.

These are not new or radical ideas when it comes to tackling the inefficiencies of the revolving door. But they are grounded in the theory that people will invest, through their favors, in cooperating at different group levels if it potentially pays to do so, and they are informed by experimental evidence.

In matters of economics and politics, group formation through favoritism is probably the primary driver of success. You can even envisage countries themselves as a loose alliance of insiders looking to do what's best for their members even when it comes at a cost to other countries. Successful countries ensure that the incentives are to cooperate at the country level and avoid costly back-scratching at lower group levels. Multi-level selection is all around.

And if you look deep enough there is an evolutionary foundation for this behavior. As Joshua Green explains:

Morality evolved to enable cooperation, but this conclusion comes with an important caveat. Biologically speaking, humans were designed for cooperation, but only with some people. Our moral brains evolved for cooperation within groups, and perhaps only within the context of personal relationships. Out moral brains did not evolve for cooperation between groups (at least not all groups). How do we know this? Because universal cooperation is inconsistent with the principles of natural selection.

Our tribalism is biological. This is why corrupt back-scratching is so hard to eradicate, and why it seems to pervade to many aspects of life.