

1

Electoral Competition under Certainty

We begin our study of formal models of domestic politics with elections. This chapter explores electoral competition when voting behavior is deterministic; the following chapter considers electoral competition under uncertainty.

Formal models of electoral competition provide a lens through which to examine the positioning of candidates in democratic elections. Consider, for example, the onetime assertion that there is little substantive difference in the platforms of the United States' two major political parties (e.g., American Political Science Association, 1950).¹ The economist Harold Hotelling was the first to offer a theoretical explanation for such behavior (Hotelling, 1929). Parties, Hotelling argued, choose positions along a left–right continuum (an example of a **policy space**), much as gas stations or drug stores choose a location along Main Street. When there are two parties, the logic of political competition compels each to adopt a position in the center of the ideological spectrum, just as we often observe gas stations located across the street from each other in the center of town. Anthony Downs popularized and extended Hotelling's argument in *An Economic Theory of Democracy* (Downs, 1957).

We thus initiate our discussion of electoral competition with the Hotelling-Downs model, where parties adopt positions to maximize their probability of winning. We then take up an alternative model in which parties are motivated not to win office for its own sake, but to achieve the best possible policy outcome. Following this, we explore electoral competition when more than two parties compete. Finally, we endogenize the number of parties (or candidates) in the election by considering various models of entry.

¹ Empirical evidence for convergence is mixed at best; see, for example, Ansolabehere, Snyder, and Stewart (2001) and Fowler and Hall (2016). In this and the following chapter we discuss some possible rationales for divergence.

1.1 The Hotelling-Downs Model

1.1.1 Euclidean Preferences

The Hotelling-Downs model is most easily expressed as a static game of complete information, where two parties simultaneously choose positions and the election outcome follows mechanically and deterministically from those policy choices.² The implicit assumption of the Hotelling-Downs model is that parties are able to credibly commit to implementing whatever policy they have promised during the election campaign. One motivation for this assumption is that parties are long-lived and therefore have an incentive to acquire a reputation for keeping campaign promises (Alesina and Spear, 1988; Cox and McCubbins, 1994; Aldrich, 1995).

We focus for now on the special case of a one-dimensional policy space, which for simplicity we assume to be the entire real number line \mathfrak{R} ; we denote any generic policy by x . In this model, there are two players, parties labeled $P = A, B$. Each party P has the same strategy space, choosing a position $x_P \in \mathfrak{R}$. Further, each party prefers outcomes that imply a higher probability of winning to those that imply a lower probability, where $\pi(x_A, x_B)$ is the probability that party A wins, given that party A and party B have chosen positions x_A and x_B , respectively.

To define $\pi(x_A, x_B)$, we describe voters' preferences and behavior and the electoral rule:

- (i) There is a continuum of voters, indexed by i , each with unique **ideal point** (most-preferred policy) $x_i \in \mathfrak{R}$. The distribution of ideal points is continuous and strictly increasing on some interval, so that there is a unique median ideal point, which we denote x_m . Voters have **Euclidean preferences** over policy, so that a voter always prefers a policy closer to her ideal point to one further away. These preferences can be represented by the utility function

$$u_i(x) = -|x - x_i|.$$

- (ii) Voters vote sincerely, choosing the party whose policy they most prefer. Voters who are indifferent between the two parties abstain.
- (iii) The election is plurality-rule: the party with the most votes wins. If the two parties receive the same vote, then the election winner is chosen by a fair lottery.

² Equivalently, we can think of the Hotelling-Downs model as an extensive game of complete information, where a finite number of voters vote strategically after parties have chosen positions. In this alternative formulation, we assume that voters play weakly undominated strategies, which, as discussed later, implies that voters in equilibrium vote for the party whose position they most prefer.

1.1 The Hotelling-Downs Model

3

Given these assumptions, $\pi(x_A, x_B)$ equals one if the fraction of voters who strictly prefer x_A to x_B is greater than one-half, equals zero if the fraction of voters who strictly prefer x_B to x_A is greater than one-half, and equals one-half otherwise.

To derive a prediction for the play of actors in this strategic environment, we look for the set of Nash equilibria. We begin by deriving the best-response correspondence for party B , that is, we find the set of optimal policy choices for party B , given x_A . Consider, for example, the optimal x_B when $x_A < x_m$. Party B can win with certainty by adopting any position closer to x_m than is x_A . (To see this, note that because voters prefer policies closer to their ideal points to those further away, party B is preferred by all voters with ideal point $x_i > \frac{x_A + x_B}{2}$, which is more than one-half of all voters given that $\frac{x_A + x_B}{2} < x_m$.) In contrast, choosing either i) x_A or ii) a position the same distance from x_m as x_A but on the other side of x_m gives a probability of winning of one-half: in (i) all voters are indifferent between party A and party B and so abstain, whereas in (ii) voters divide evenly between party A and party B . Finally, choosing a position further away from x_m than is x_A means that party B loses with certainty.

A similar logic applies when $x_A > x_m$. Thus, when $x_A \neq x_m$, any position closer to x_m than is x_A is a best response. Finally, when $x_A = x_m$, only x_m is a best response: choosing $x_B = x_m$ results in a probability of winning of one-half, whereas any other position entails losing with certainty.

Party A 's best-response correspondence is analogous: if $x_B \neq x_m$, any position closer to x_m than is x_B is a best response, whereas if $x_B = x_m$, the best response is x_m . Each party's best response is therefore to choose a position closer to x_m than is the other party's position, when that is possible. Clearly, the two parties are playing a best response to each other only when $x_A = x_B = x_m$. This is the unique Nash equilibrium. The logic of political competition forces each party to adopt the median ideal point, as only when that is the case is neither party able to increase its probability of winning.

1.1.2 Single-peaked Preferences

The assumption that voters have Euclidean preferences, though convenient, is restrictive. In many policy environments, it is natural for voters to have asymmetric preferences, valuing differences to one side of their ideal point more than those to the other. We should therefore ask whether the result obtained in the previous section—that in equilibrium parties each adopt the median ideal point—carries through if we assume more generally that voters have single-peaked preferences, which we define as follows:

Voters have **single-peaked preferences** over policies in \mathfrak{R} if and only if, for each voter i , there is a unique ideal point x_i and the following condition holds for all $x', x'' \in \mathfrak{R}$:

$$\text{if } x'' < x' < x_i \text{ or } x'' > x' > x_i, \text{ then } x' \succ_i x'',$$

where \succ is the strict preference relation.

Preferences are single-peaked with respect to policies along the real number line if and only if each voter has a unique ideal point and—among positions *on the same side* of that ideal point—prefers positions that are closer to the ideal point to those further away. Clearly, Euclidean preferences are a special case of single-peaked preferences.

Social choice theory tells us that if individuals have single-peaked preferences, then an alternative is a **Condorcet winner** (an alternative such that no other alternative is strictly preferred by a majority) if and only if it is a median ideal point. The same logic implies that if voters in a Hotelling-Downs environment have single-peaked preferences, then the parties converge to a median ideal point, as any other position can be beaten.

To see this, assume as before that the distribution of voters' ideal points is continuous and strictly increasing on some interval, so that there is a unique median ideal point. Our intuition is that (x_m, x_m) is the unique Nash equilibrium, that is, that this strategy profile is a Nash equilibrium and no others are. We first demonstrate existence ((x_m, x_m) is a Nash equilibrium) and then uniqueness (no other strategy profile is a Nash equilibrium).

- (i) Existence: When the parties adopt (x_m, x_m) , each party wins with probability one-half. Without loss of generality (that is, the same argument applies to party B), consider a deviation by party A to some $x' < x_m$. As party B is preferred by all voters with ideal point to the right of x_m , party B wins with probability at least one-half.³ Thus, there is no profitable deviation.
- (ii) Uniqueness: We prove that (x_m, x_m) is the unique Nash equilibrium by showing that for all other strategy profiles at least one party has an incentive to deviate. Consider three mutually exclusive and exhaustive cases:
 - (a) *One of the parties wins with certainty.* The losing party can adopt the position chosen by the winning party and win with probability one-half. Thus, this is not a Nash equilibrium.

³ In fact, party B wins with certainty, though this is harder to show and not essential to the proof. See Roemer (2001, Section 1.2).

- (b) *Parties A and B each win with probability one-half, with $x_A = x_B \neq x_m$.* Without loss of generality, assume $x_A = x_B < x_m$, and consider a deviation by party B to some x' such that $x_A < x' < x_m$. With this deviation, party B is preferred by all voters with ideal point to the right of x' and wins with certainty. Thus, this is not a Nash equilibrium.
- (c) *Parties A and B each win with probability one-half, with $x_A < x_m < x_B$ or $x_B < x_m < x_A$.* Without loss of generality, assume $x_A < x_m < x_B$. Then by the argument for the previous case, party B can profitably deviate to some x' such that $x_A < x' < x_m$. Thus, this is not a Nash equilibrium.

The logic of the proof illustrates another insight of the Hotelling-Downs model: two-party elections are often close. So long as parties have the freedom to commit to any position in the policy space, either party can guarantee itself a tie by adopting the position chosen by the other. In equilibrium, therefore, each party wins with probability one-half.

1.1.3 Hotelling-Downs Competition in a Multidimensional Policy Space

We are conditioned to think of politics as one-dimensional. Politicians and journalists speak of “liberal” and “conservative” policies, and parties throughout the world are labeled “leftist” or “rightist.” Yet even simple policy environments may be inherently multidimensional. Consider, for example, the “pie-splitting” environment, where three individuals must decide how to divide a “pie” of size 1. Let q_1 be the share received by individual 1 and q_2 that received by individual 2, so that individual 3 receives $1 - q_1 - q_2$. Assume that individuals prefer more pie to less.

There is no Condorcet winner in this environment. To see this, assume to the contrary that a Condorcet winner (q_1, q_2) exists. Because the shares of the three individuals sum to one, it must be true that at least one individual receives something from this policy. But then that individual’s share could be divided between the remaining two players, who would clearly prefer this alternative division (q'_1, q'_2) to (q_1, q_2) .

A similar logic applies in the context of electoral competition. Consider the Hotelling-Downs model, but now assume that the parties compete by proposing a division of a “pie” of size one among three groups, labeled $g = 1, 2, 3$. Let α_g be the size of group g , with $\sum \alpha_g = 1$ and $\alpha_g < \frac{1}{2}$ for

all g ; thus, any two groups constitute a majority. Again, individuals prefer more pie to less. We denote by (q_{1P}, q_{2P}) the policy offered by party P .

There is no Nash equilibrium of this game.⁴ To see this, assume to the contrary that there is some strategy pair $((q_{1A}, q_{2A}), (q_{1B}, q_{2B}))$ that is a Nash equilibrium. Note that in this equilibrium either one party wins with certainty or the two parties each win with probability one-half. In the first case, the losing party can increase its probability of winning to one-half by choosing the same policy as that chosen by the other party. In the second case, either party can increase its probability of winning to one by adopting a policy preferred by two groups to the policy chosen by the other party; by the argument already given, such a policy exists. Thus, there is no Nash equilibrium.

Intuitively, when there is no Condorcet winner, then there is no equilibrium to the Hotelling-Downs model, as any policy can be beaten by some other policy. As we show in the following chapter, however, this result is sensitive to the assumption that individuals' voting decisions follow deterministically from their policy preferences.

1.2 The Wittman Model

Up to now, we have assumed that parties care only about winning. We might defend this assumption by arguing that the nonpolicy benefits of holding office (prestige, patronage power, etc.) are paramount. The universality of this argument, however, is questionable: many politicians appear to enter politics not for the perks of office but because of their strong policy preferences. It is intuitive that parties made up of such politicians would be less likely to compromise on policy for the sake of winning office.

Donald Wittman was the first to formulate a model with **policy-seeking** rather than **office-seeking** parties (Wittman, 1973).⁵ Surprisingly, our intuition that policy-seeking parties may be less inclined to adopt centrist positions does not hold in the basic Wittman model: as in the Hotelling-Downs model, the unique equilibrium is for each party to adopt the median ideal point. Intuitively, even though parties care about policy, they can implement policy only by winning office. The logic of political competition therefore drives them to adopt the same centrist policies they would choose if they were instead motivated to win office for its own sake.

⁴ Here and in subsequent sections we implicitly restrict attention to pure-strategy Nash equilibria. In general, discontinuities in the parties' payoff functions in the Hotelling-Downs and related models imply that mixed-strategy Nash equilibria may not exist.

⁵ In general, the literature assumes that parties are either all office-seeking or all policy-seeking. An exception is Callander (2008a), who considers heterogeneous motivations.

To focus the discussion, assume as in the model of Section 1.1.1 that voters have Euclidean preferences over $x \in \mathfrak{R}$. There are two parties, $P = L, R$, which have von Neumann-Morgenstern preferences over lotteries over policy outcomes, where L receives a payoff equal to $-|x|$ and R a payoff equal to $-|x - 1|$ if policy x is implemented (i.e., the parties have ideal points 0 and 1, respectively). We assume $0 < x_m < 1$, so that the parties are “polarized.” Parties L and R maximize expected utility by choosing positions x_L and x_R , respectively. Thus, letting $\pi(x_L, x_R)$ be the probability that L wins, given (x_L, x_R) , L solves

$$\max_{x_L} \pi(x_L, x_R)(-|x_L|) + [1 - \pi(x_L, x_R)](-|x_R|),$$

whereas R solves

$$\max_{x_R} \pi(x_L, x_R)(-|x_L - 1|) + [1 - \pi(x_L, x_R)](-|x_R - 1|).$$

There is a Nash equilibrium of this game in which each party chooses x_m . Proving this is easy: if either party deviates to some other position, then that party loses with certainty rather than winning with probability one-half. Because losing to a party that has adopted x_m gives the same expected utility as winning with probability one-half when each party has adopted x_m , there is no profitable deviation.

Moreover, this is the unique Nash equilibrium, though showing that involves a few more steps. The basic logic can be seen by assuming that the parties have chosen positions $0 < x_L < x_m < x_R < 1$, with $|x_L - x_m| = |x_m - x_R|$. Because the median voter is indifferent between the two parties, each party wins with probability one-half. Thus, for example, L has expected utility $-\frac{1}{2}(x_L + x_R) = -x_m$. However, L can profitably deviate by moving some infinitesimal ϵ to the right, increasing its probability of winning from one-half to one, and thus receiving expected utility $-(x_L + \epsilon) > -x_m$. Intuitively, divergence is not a Nash equilibrium, as there is always an incentive to move a bit closer to the center and thus win for sure. As we will see in the next chapter, this incentive is softened when policy preferences map stochastically onto voting decisions, as then a small move toward the center results in only a small increase in the probability of winning.

1.3 Multiparty Competition

Our discussion so far has been limited to models of two-party competition. Such models were the focus of most early formal work on electoral competition, perhaps due to the predominance of two-party competition in the

United States. However, in many political environments, more than two parties compete for the vote. Following the literature, we refer to this as multiparty competition, though a literal interpretation of this term would also include two-party competition.

As a point of departure, consider the Hotelling-Downs model with one-dimensional policy competition, but now assume that the election is contested by three parties, $P = A, B, C$, each of which maximizes its probability of winning. We continue to assume that voters vote mechanically for the party they most prefer, though this assumption is far less innocuous in a multiparty setting. We take up the question of strategic voting later this chapter. Further, we adapt the model of voter behavior from that considered earlier by assuming that if voters are indifferent among two or more parties, then they choose a party from among those they most prefer using an equal-probability rule (e.g., they flip a fair coin if they are indifferent between two parties) rather than abstaining.

One might expect that each party would adopt the median ideal point, as with two-party competition. However, this is not a Nash equilibrium: any party could profitably deviate by adopting a position some arbitrarily small distance from x_m and thus receiving almost half the vote, leaving its two competitors to divide the remaining half. As this is a plurality-rule election, any party that deviated in this way would win with certainty.

So is there a Nash equilibrium of this game? Yes, for certain distributions of voter preferences. Assume, for example, that voters have Euclidean preferences with ideal points distributed uniformly on $[0, 1]$. Then the following configuration of positions is a Nash equilibrium:

$$\begin{aligned} x_A = x_B &= \frac{1}{3}, \\ x_C &= \frac{2}{3}. \end{aligned}$$

To see this, note that in equilibrium party C receives one-half of the vote and wins with certainty, so that party C has no incentive to deviate. Consider possible deviations by party A ; the same arguments apply to party B :

- Adopting any $x_A < \frac{1}{3}$ leaves party C with one-half of the vote while dividing the other half between parties A and B , so that party C continues to win with certainty.
- Adopting some $x_A \in (\frac{1}{3}, \frac{2}{3})$ provides one-sixth of the vote and divides the remaining five-sixths between parties B and C . Thus either party B or C wins with certainty or, if $x_A = \frac{1}{2}$, parties B and C each win with probability one-half.

1.3 Multiparty Competition

9

- Adopting $x_A = \frac{2}{3}$ divides one-half of the vote between parties A and C , so that party B wins with certainty.
- Adopting $x_A > \frac{2}{3}$ provides party B with one-half of the vote while dividing the other half between parties A and C , so that party B wins with certainty.

The equilibrium in this example is nonetheless unattractive as an empirical prediction. If both parties A and B expect to lose, then we may ask why they entered the race to begin with. The answer may be that the fixed cost of competing in this particular electoral arena was previously sunk. In that case, parties may choose to contest elections even when they expect to lose so that they survive to contest future elections, or so that they may contest elections in other arenas. We turn to the question of endogenous entry in elections later in this chapter.

Further, we may ask why party C is content merely to win, when moving toward $x = \frac{1}{3}$ would increase party C 's vote share. Clearly, the configuration of positions above is not a Nash equilibrium when parties maximize vote share. More generally, there is no Nash equilibrium of the game with these alternative preferences, even when voters' ideal points are uniformly distributed. To see this, consider the following list of conditions, which Cox (1987b) establishes are necessary for any Nash equilibrium when parties maximize vote share:⁶

- (i) No more than two parties occupy any one position.
- (ii) Each extremist position (meaning a position leftmost or rightmost among those occupied by the parties, not leftmost or rightmost among all positions that could be occupied) is occupied by exactly two parties.
- (iii) If two parties occupy the same position x , then the share of voters to the left of x who most prefer x among all positions that have been adopted is equal to the share of voters to the right of x who most prefer x among all positions that have been adopted.

With three parties, Condition (ii) cannot be satisfied.

What is the appropriate objective for office-seeking parties in multiparty competition? In two-party competition maximizing (expected) vote share and maximizing the probability of winning are often equivalent. However, as the example just given suggests, this is not generally the case in multiparty competition. Choosing the appropriate objective therefore boils down to whether one believes that a larger vote share translates into additional

⁶ These conditions hold not only when parties maximize vote share, but also when they maximize the margin of victory relative to the second-place finisher and when they maximize their "total" margin of victory relative to all other parties.

post-election benefits. In most environments, it seems that it must. In parliamentary systems, for example, a large vote share may increase the probability that a party controls the policy agenda, a consideration we take up when examining models of legislative bargaining in Chapter 6. For the remainder of this section we therefore assume that parties maximize vote share rather than their probability of winning.

The conditions just listed ensure that there is no equilibrium in three-party competition when parties maximize vote share. What about four-party competition? For certain distributions of voter preferences, such an equilibrium may exist, with parties adopting divergent positions. Assume, for example, that citizens have Euclidean preferences and ideal points distributed uniformly on $[0, 1]$. Then the following configuration of positions is a Nash equilibrium of the Hotelling-Downs model with four-party competition:

$$\begin{aligned}x_A = x_B &= \frac{1}{4}, \\x_C = x_D &= \frac{3}{4}.\end{aligned}$$

To see this, note that in equilibrium every party receives one-fourth of the vote. Without loss of generality, focus on possible deviations by party A . Adopting any $x_A < \frac{1}{4}$ or $x_A > \frac{3}{4}$ provides less than one-fourth of the vote; adopting any $x_A \in (\frac{1}{4}, \frac{3}{4})$ gives exactly one-fourth of the vote; and adopting $x_A = \frac{3}{4}$ provides one-sixth of the vote. Thus, party A has no incentive to deviate.

An equilibrium exists in this example because the parties are spaced in such a way that no party can increase its vote share by adopting a position arbitrarily close to that of another party. However, equilibrium existence is sensitive to the distribution of voters' ideal points, as demonstrated by another example. Assume that citizens have Euclidean preferences and ideal points x_i distributed on $[0, 1]$ according to a triangular distribution with distribution function $F(x_i) = 2x_i - x_i^2$. The first two conditions given earlier imply that the parties must be paired, with two parties occupying one position x_L and the other two parties occupying some other position $x_R > x_L$. Let $\hat{x} = \frac{x_L + x_R}{2}$ be the ideal point of voters who are indifferent between the positions adopted by the four parties. Then the third condition implies that $F(x_L) = F(\hat{x}) - F(x_L)$, and $F(x_R) - F(\hat{x}) = 1 - F(x_R)$.

Figure 1.1 shows the necessary relationships among different subpopulations: the fraction of voters to the left of x_L (which we denote by a) must equal that between x_L and \hat{x} (which we denote by b), and the fraction of voters between \hat{x} and x_R (which we denote by c) must equal that between