# Chapter 8
# Economic Models: Gods, Supplicants, and Priests

**Abstract** This chapter provides four formal models to match the previous economic analysis. First, a model of divine preferences constructed as a sequence of Edgeworth boxes, which illustrates the supplicant's dilemma and the theological escapes from it. Second, a model of a Prisoner's Dilemma game among priests, illustrating the incentive to defect from cooperation and the ways to overcome it. Third, a model of an Assurance game between gods and priests which has two institutional equilibria, a risk-dominant (Greco-Roman) equilibrium and a payoff-dominant (Hindu and Zoroastrian) equilibrium. Fourth, a model of the missionary expansion of a cooperative religious organization that protects the old members' benefits as a condition for the expansion to be acceptable, illustrating the Hindu and Zoroastrian expansions.

## 8.1 A Model of Divine Preferences

### *8.1.1 The Supplicant's Dilemma*

For ease of exposition, we start with the benchmark situation of overlap of divine jurisdictions, which describes the archetypal problem that all polytheistic systems had at least potentially to cope with.[1] Then we will depict the different ways in which the problem could be avoided, overcome, or otherwise kept under control.

Imagine that there are several gods who have overlapping, but non-coincident, jurisdictions over several matters. (If the jurisdictions were wholly coincident, the gods would be perceived as identical, which would raise questions about their rivalry—see below.) As so often in economics, a $2 \times 2$ model will suffice to illustrate the problem. Let A and B be two gods who are thought to be able to affect outcomes in two fields, x and y. For example, x might be a woman's health and y her giving birth, and A and B might be Hera and Artemis; or, x might be victory in war and y wealth, and A might be Iran's Mithra and B Indra—both war gods and both gods of material plenty, the first as the lord of covenant and justice and the second through

[1] This section is based on the models developed in Ferrero and Tridimas (2018), Basuchoudhary, Ferrero and Lubin (2020) and Ferrero (2021).

 
M. Ferrero, *The Political Economy of Indo-European Polytheism*,
Contributions to Economics, https://doi.org/10.1007/978-3-030-97943-0_8

plunder and war booty. Let $x_A$ and $x_B$ (respectively, $y_A$ and $y_B$) denote the amounts of sacrifices or offerings of an agent to gods A and B in pursuit of a favorable outcome in field x (respectively, y). Suppose for the moment that there are fixed amounts of resources, $\overline{x}$ and $\overline{y}$, that the agent can or will devote as offerings to each outcome, such that $x_A + x_B \leq \overline{x}$ and $y_A + y_B \leq \overline{y}$.

The agent maximizes utility on behalf of each god by choosing the offering basket with a view to satisfying the god as best he can. Analytically, gods are perceived to be pleased by the offerings dedicated to them, thankful for them, and willing to reward the supplicant by bestowing favors on him towards fulfillment of the supplicant's wishes. These wishes, and the corresponding benefits expected from the god, can be worldly and/or otherworldly. The happier the god, the fuller and more effective are his/her blessings—which makes ancient polytheism a kind of "fee for service" operation (Iannaccone 1995), although one founded on subjective beliefs and subject to random disturbances. However, the supplicant is well aware that each god is sensitive to his/her being recognized as influential in both fields and will not make the mistake of only trusting one god for $x$ and the other for $y$—unless, that is, the gods are perceived as specializing in $x$ or $y$ (see below).

Furthermore, each god is supposed to have a satiation point which is, in principle, within the supplicant's reach and consists of a bundle of offerings that makes the god perfectly satisfied. This point is again a subjective belief of a typical supplicant, as determined by the prevailing culture of the ancient world which viewed the gods as eager for acknowledgment and offerings but amenable to be pacified with sufficient effort. The theology, mythology, and cult practices of ancient polytheistic religions strongly suggest that such satiation points were thought to be knowable by the supplicants and/or the priests—they "knew" what the gods wanted—and relatively immune to disconfirmation from perceived failures of past offerings to fulfill one's wishes—that is, they were thought to be stable. This is the confirmation bias (or asymmetric valuation of errors) discussed in the previous chapter. Hence, a dynamic model with updating of beliefs and endogenous satiation points does not seem necessary to handle the problem. People behaved as if they felt that sometimes their offerings were adequate, sometimes not, that is, sometimes their wishes were fulfilled, sometimes not; and they tried hard to minimize the danger of the gods' displeasure.

If the supplicant cannot or will not make offerings that match the satiation point, he believes he will face in return less satisfactory or more haphazard blessings from the god. Any offering in excess of this ideal bundle, however, does not turn the "good" into a "bad" but into a "neutral" good—i.e., the god is indifferent to the excess offerings which, therefore, would neither benefit nor harm the supplicant. Hence, the supplicant will not waste scarce resources in excess of satiation levels but will turn them over to the other god or, if the resources are sufficiently flexible in use, to the other outcome.

This suggests that the resources devoted to each outcome can be made variable subject to an overall resource constraint. Suppose that $x$ and $y$ are measured in the same units—such as money or time devoted to religious observance—so that they are perfectly substitutable across outcomes, and denote with $R$ the fixed amount[2] of total resources available for religious offerings, with $R = \overline{x} + \overline{y}$. Let $S_A$ and $S_B$ be the two gods' satiation bundles and the superscript $S$ attached to the $x$s and $y$s denote the corresponding satiation levels. Avoiding waste of religious resources requires:

$$\overline{x} \leq x_A^S + x_B^S \; and \; \overline{y} \leq y_A^S + y_B^S \tag{8.1}$$

Scarcity of resources will prevail if at least one of the inequalities in (8.1) holds as a strict inequality ("<"), i.e. the supplicant cannot supply both gods with their satiation bundles. Under scarcity, to maximize utility, the available amounts of both religious resources will be fully used up, hence $x_A + x_B = \overline{x}$ and $y_A + y_B = \overline{y}$. Therefore under scarcity, using (8.1), the following inequality holds:

$$R = \overline{x} + \overline{y} = x_A + x_B + y_A + y_B < x_A^S + x_B^S + y_A^S + y_B^S \tag{8.2}$$

Scarcity of resources is, however, not sufficient to create a nontrivial choice problem for the supplicant. Two cases can be distinguished. The first is the case of *specialized gods*, where each god has exclusive jurisdiction over one field and cares only about offerings in that field, so there is no jurisdictional overlap. In this case two of the satiation levels in (8.1) and (8.2), one for each god, would be equal to zero. For example if god A cares only about $y$ and god B cares only about $x$, then $x_A^S = 0$ and $y_B^S = 0$. Then there is no real choice to make: the supplicant will naturally devote the whole of offering $y$ to god A and the whole of offering $x$ to god B, although falling short of fully satiating them if he labors under scarcity. The second is the case of *rival gods*, where each god has jurisdiction over both issues and cares about both offerings, implying jurisdictional overlap—which is where we started above. Joined with scarcity, this creates a meaningful choice problem for the supplicant.

With these assumptions, the situation can be simply modeled by an unusual application of a straightforward tool: an Edgeworth box, which depicts the indifference maps over the perceived preferences of gods A and B with respect to offerings $x$ and $y$. The sides of the box are measured by any pair $(\overline{x}, \overline{y})$ that satisfies (8.1), and which also satisfies (8.2) if there is scarcity. All the Edgeworth boxes introduced below, each corresponding to a different theology, are rectangles with exactly the same sides, allowing us to compare the allocations of resources and to see the savings or Pareto-improvements that certain theological reforms make possible. Scarcity in both dimensions implies that the inequalities in (8.1) both hold as strict inequalities, i.e. the two satiation points cannot both be reached with the existing resources; on the other hand, our assumption that each god's satiation is within the supplicant's reach implies that both satiation points lie within the box (including its edges). Hence it must be $\overline{x} = max\left(x_A^S, x_B^S\right)$ and $\overline{y} = max\left(y_A^S, y_B^S\right)$, i.e. the length of each side of

[2] A full model would determine, in the usual way, the supplicant's optimal $R$ by maximizing his utility over religious and secular consumption subject to a budget constraint. At an interior solution, one would expect this $R$ to fall short of the gods' satiation bundles, i.e. scarcity would prevail. Giving priority to satiation of both gods regardless of opportunity cost in terms of forgone consumption would seem to require some kind of lexicographic preferences.

**Fig. 8.1** Rival gods: both A and B care about both *x* and *y*

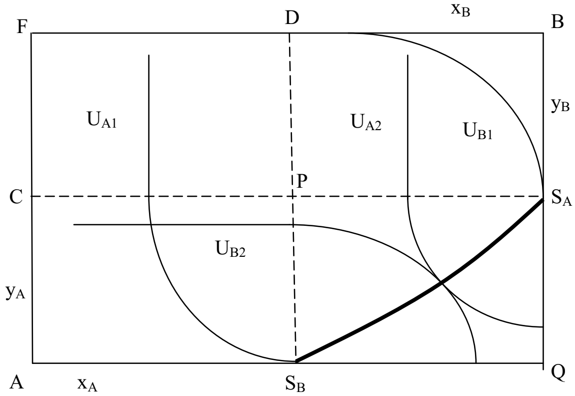


the rectangle must be exactly equal to the highest of the two satiation levels that are measured on that side, no more (or there would be offerings exceeding the satiation levels, implying waste of resources) and no less (or the satiation points would fall outside the box). It follows that both $S_A$ and $S_B$ will lie somewhere on the sides of the box. If the two gods were identical, with wholly coincident jurisdictions ($x_A^S = x_B^S$ and $y_A^S = y_B^S$), the satiation point of each god would coincide with the origin vertex for the other god, i.e. fully satiating one god would leave the other god with a bundle $(x, y) = (0, 0)$. We leave this special case for further discussion below and focus on the general case of non-identical gods, where the two satiation points lie on two adjacent sides of the box. For concreteness, and without loss of generality, we set $x_A^S = 2\, x_B^S$ and $y_B^S = 2y_A^S$, so in our boxes below $\overline{x} = x_A^S$ and $\overline{y} = y_B^S$. Hence, here total expenditure on religious offerings is $R = x_A^S + y_B^S$.

With these assumptions, Fig. 8.1 depicts our benchmark case of *rival gods*, where each god is perceived to have jurisdiction over both issues and the two gods' indifference maps culminate at satiation points $S_A$ and $S_B$. The pair of horizontal and vertical dashed lines drawn through each satiation point divide the $(x, y)$ space into an area where both goods are below satiation levels and indifference curves are monotonically increasing in utility and strictly convex to the origin, and the rest of the space where one or both goods are above satiation levels and indifference curves become straight lines because the excess offering leaves the god indifferent—i.e. the good becomes a "neutral". Two such indifference curves are drawn for each god. The SE quadrant $S_A P\, S_B Q$ is the only subset of the box where both goods are below satiation levels for both gods, and it is easy to check that a move from any point outside $S_A P\, S_B Q$ to a point inside it represents a Pareto improvement whereby one or both gods increase their utility. The thickened curve $S_A S_B$ connects indifference curves $U_{A1}$ and $U_{B1}$ (A's indifference curve through B's satiation point and B's indifference curve through A's satiation point) and is the locus of tangency points between pairs of indifference curves; hence, a move from a point anywhere in the rectangle to a point on the curve represents a Pareto improvement. Thus $S_A S_B$ is the "contract curve" between the two gods, i.e. the set of Pareto optimal allocations of offerings between

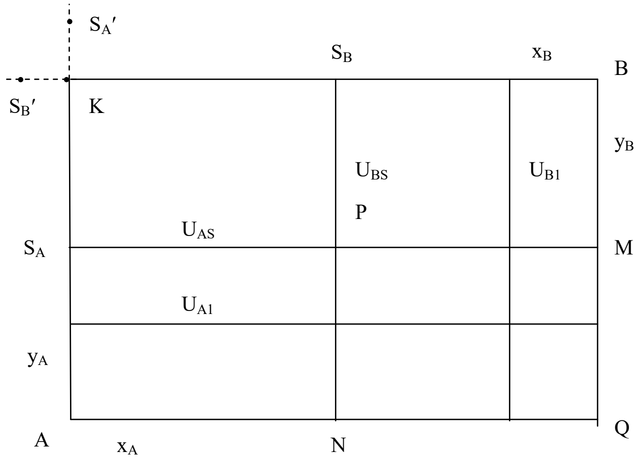


**Fig. 8.2** Specialized gods: A cares only about *y*, B cares only about *x*

them. It depicts the trade-off faced by the supplicant: by increasing his offerings to placate one god, he will incur the displeasure of the other god. Along the contract curve, the gods are indeed "jealous" of each other as any change in allocation that makes one god better off will make the other worse off. Choosing an allocation on the $S_A S_B$ curve, not outside of it, is thus the best that can be done under the existing resource constraint to minimize the harm from gods' displeasure.[3] The "hungrier", or more demanding, are the gods, given the resources, the further apart are their satiation points and the worse is the trade-off.[4] This is the *supplicant's dilemma*.

This is to be contrasted with the case of *specialized gods*, in which there is no jurisdictional overlap. There are two possible patterns of specialization, one of which is depicted by Fig. 8.2, with god A influencing only and therefore caring only about *y* and god B caring only about *x*. Here their indifference curves become everywhere horizontal and vertical lines, respectively, starting from points on the sides of the box, because the good about which the god does not care is a neutral; the corresponding

[3] The contract curve implicitly defines a "utility possibility frontier" of the gods, where the supplicant will want to choose his preferred point based on his beliefs and attitudes to risk. In a standard case of risk aversion, he will typically pick an interior solution, thus diversifying his portfolio of offerings and spreading the risk. Modeling this would be straightforward but would add little of interest for our purposes.

[4] In theory, there could be "bargaining between the gods" to find a mutually acceptable sharing of the offerings, i.e. a particular point on the contract curve, for example a Nash solution to the bargaining problem. One might perhaps interpret the elaborate system of festivals and ceremonies fixed by the Roman state as embedding some such solution—a risk-minimizing arrangement. Elsewhere, this did not happen. The reason is probably to be sought in the fact that either (in India, Iran, and the Celtic world) the priesthood was not specialized by god, so the different gods were not "represented" by different priests who would bargain on their behalf, or (in Greece) it was hyper-specialized by local temples of local specifications of a god, so again unable to represent "the" god even if they had sufficient incentives to do so (which they did not—see Sect. 8.2 below).

utility levels, including satiation levels, are measured by the quantities of the only good that the god cares about, denoted by points on the sides of the box. If scarcity prevails for both resources so that both satiation points are beyond the supplicant's resources and hence outside the box, like $S_A{}'$ and $S_B{}'$, he will choose the NW vertex of the box (point $K$) as his unique optimal bundle, exhausting his available resources. If, however, the satiation points are $S_A$ and $S_B$—corresponding to the satiation levels of the relevant goods in the benchmark satiation points of Fig. 8.1—the supplicant would be in a state of bliss: both gods could be satiated without using up all the available resources. In either case, no "trade" of offerings between the gods is required or feasible: each god receives all and only the available amount of the good he cares about. There is no supplicant's dilemma. Of course, in the latter case an excess of resources is being needlessly tied down to religious offerings, violating (8.1); to maximize his utility, the supplicant should reduce these to the benefit of secular consumption and downsize the box to $PMQN$, where $S_A$ and $S_B$ would be made to coincide with point $P$, which would be the unique optimum.

Consider, however, the opposite specialization pattern (not shown), with god A specializing in good x and god B specializing in good y, and with the satiation level of each good for each god again corresponding to that of the relevant satiation point of Fig. 8.1. All utility levels, including satiation levels, are still measured by points on the sides of the box and all indifference curves inside the box are still straight lines, but the two satiation points now coincide at the SE vertex $Q$, which is the unique optimum. This optimum achieves full satiation of both gods, so again there is no trade-off and no supplicant's dilemma, but here no downsizing of the box and thus no saving of religious resources is possible because it would put the satiation points out of reach. Why the difference? Recall that we assumed $x_A^S > x_B^S$ and $y_B^S > y_A^S$; so in the pattern depicted by Fig. 8.2 each god specializes in the good for which his satiation level is lower than that of the other god, and therefore is "cheaper" to satisfy, whereas in the alternative pattern it is higher than the other's and therefore more expensive to satisfy. That explains why saving of resources, compared to the rival gods situation, becomes feasible in the former case but not in the latter. In both cases, however, specialization gets rid of gods' jealousy and the associated supplicant's dilemma.

The specialized gods model, with no jurisdictional overlap, no divine jealousy, and no dilemma, may capture the essence of some of the earliest Indo-European pantheons on record: in particular those of the Thracians described by Herodotus, and possibly the earliest Gauls and Germans as described by the Roman historians Caesar and Tacitus (see Sects. 1.5, 3.1.1 and 3.2.1). In these cases, the specialization has been claimed to conform to Dumézil's tri-functional structure. The model may also be used to describe the relationship between the various *yazatas* at the lower levels of the Zoroastrian hierarchy, where formerly independent gods have been turned by Zoroaster into subordinate, specialized divinities, each appointed to a particular, non-overlapping domain or function. In fact, Dumézil (1958, 40–46) claimed that despite—or rather, thanks to—Zoroaster's reform, the list of the Amesha Spentas in particular mirrors the classical tripartite structure of the earlier Indo-Iranian and Vedic pantheon. At about the same time as the Thracians, however, Herodotus'

Scythians hardly conformed to any well-defined specialization pattern, tri-functional or otherwise. Moreover, as we have seen (Chap. 3), the later-documented brethren of the Gauls and the Germans (the Irish and the Vikings respectively) show evidence of multiplication of deities and/or of shifting and broadening specializations, suggesting that the specialized structure was fragile and ill-equipped to stand the tests of time, migration, and social change.

Instead, the rival gods model seems apt to capture these later evolutions, and in general to describe the central problem that, in different forms and degrees, was at least potentially undermining all the Indo-European pantheons discussed in this book. This problem was described in the last chapter as the increase in jurisdictional overlap, brought about by the proliferation of deities and/or the broadening of their original specializations; this increase, if unchecked, was likely to harden divine jealousies and consequently to increase the burden of cult and the attendant cost and anxiety borne by supplicants in an effort to cope with them. As we have seen, this description, and hence the model, fits particularly well the early Greek and Roman pantheons and even more so the unified Greco-Roman pantheon that capped them, where the number-overlap-jealousy complex rose to unparalleled proportions. We will now develop different specifications of the Edgeworth box model to capture the various ways around the problem and out of the dilemma taken by some of the Indo-European peoples.

### *8.1.2 Escapes from the Dilemma*

The Hindu response to the jealous gods problem was theistic sectarianism. Figure 8.3 depicts the theology of sectarian Hinduism, with the same box size and the same satiation points as in Fig. 8.1. As before, resources are scarce as the supplicant

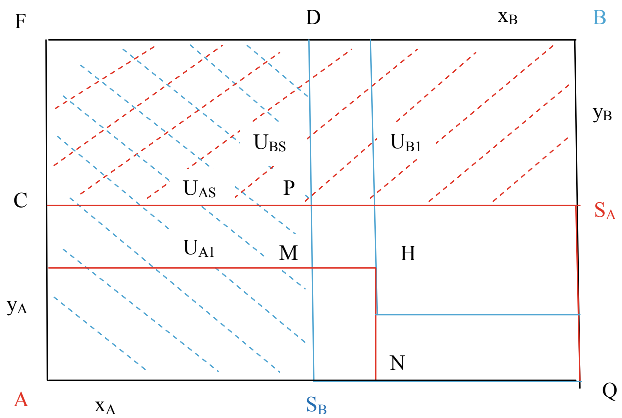


**Fig. 8.3** Sectarian Hinduism: non-jealous gods

cannot simultaneously provide the gods with the bundles $S_A$ and $S_B$. Here, however, the gods' preferences are monotonic but strictly *concave*, meaning that "extremes are preferred to averages", i.e., the agent prefers to specialize, at least to some degree, and to consume only or mostly one good rather than a mixed bag. For ease of illustration and with little loss of generality, in this and the following figures, we consider the limit case of indifference curves that are inverted L-shaped, i.e. with the vertex of the L placed *opposite* to the origin. For clarity, all the points and curves for A are drawn in red, for B in blue; we show two indifference curves for each god. These curves are derived from utility functions $U_A = max(x_A, \alpha y_A)$ and $U_B = max(x_B, \beta y_B)$, in which $\alpha$ (respectively, $\beta$) is the constant offerings ratio $x_A/y_A$ (respectively, $x_B/y_B$) at the vertex of all indifference curves, where the two arguments take the same value which measures the utility level. This means that god A's satisfaction is measured by the amount of offering that is the *greatest* of $x_A$ and $\alpha y_A$; given this, the amount of the other offering is irrelevant (and similarly for god B). For example, consider bundle $H$ from A's point of view (red line): if we move horizontally to $M$ we decrease $x_A$ while leaving $y_A$ unchanged, so utility does not decrease because the greater amount (now $\alpha y_A$) is unchanged; hence $H$ and $M$ belong to the same indifference curve. While at a cosmic level of abstraction, this captures the idea that Shiva or Vishnu value and appreciate the supplicant's effort and intention to please and acknowledge them despite his limited means—that is, his devotion *(bhakti)* rather than the amount of stuff that they themselves "consume."

To find the efficient solutions, start from $H$: this is not an optimal point as moving to $M$ (or $N$) leaves A indifferent but increases B's utility—indeed B has now a bundle that lies on his indifference curve ($U_{BS}$) through $S_B$, i.e., it gives him the same utility as satiation. But $M$ (or $N$) is not optimal either as moving from there to $P$ (or $Q$) leaves B indifferent (at satiation) and increases A's utility to the level of satiation as $P$ (or $Q$) lies on his indifference curve ($U_{AS}$) through $S_A$. So $P$ and $Q$ are Pareto-optimal allocations: the remarkable feature is that at these points, due to the gods' concave preferences, resources are still scarce, but divine jealousy is no more; that is, the gods are not rivalrous.

Furthermore, since we assumed that offerings in excess of a god's satiation level neither please nor displease him (they become "neutrals"), all the bundles in the area marked with red (blue) dashed lines, which lie above the indifference curve through $S_A$ ($S_B$), are indifferent to $S_A$ ($S_B$). Therefore the set of Pareto-optimal allocations comprises the rectangle to the NW of $P$ where the red and blue dashed lines overlap, up to the vertex of the Edgeworth box $F$, and at any of those allocations, both gods are as happy as at their respective satiation points. For example, at point $P$ god A is receiving his satiation level of $y$ but not of $x$, yet this is just as good to him as $S_A$, and god B is receiving his satiation level of $x$ but not of $y$, and yet this is just as good to him as $S_B$. Then, if desired (for example, to economize on transaction costs), the supplicant can just as well specialize his offerings entirely and go to point $F$, thus giving all of his $y$ offerings to A and all of his $x$ offerings to B; Shiva and Vishnu will not mind. Thus, with these non-jealous gods, the supplicant's dilemma goes away.

To note that $P$ and $F$ are equivalent optima is to imply that valuable resources are being unnecessarily tied down in excess religious offerings, so even if an allocation within this box may be efficient, the box itself is oversized and can and should be downsized to everyone's satisfaction. Suppose that both $\overline{x}$ and $\overline{y}$ are reduced and the box is squeezed in both dimensions from above and from the left, so that the NW vertex $F$ is made to coincide with point $P$. Figure 8.4 represents the new box, where points $S_A$, $P$, $S_B$, and $Q$ are reproduced from Fig. 8.3 for ease of comparison. God A's origin is now $A'$, coincident with old $S_B$, and god B's is now $B'$, coincident with old $S_A$, which shifts their satiation points to $S_A{}'$ and $S_B{}'$ respectively (segments $CS_A$ and $DS_B$ in Fig. 8.3 are equal to segments $PS_A{}'$ and $PS_B{}'$ in Fig. 8.4, respectively). These new satiation points lie outside the box, but the new indifference curves running through them ($U_{AS}{}'$ and $U_{BS}{}'$ respectively) still cross at point $P$, which is now the unique optimum. Clearly, no further squeezing of the box is possible because it would put satiation out of reach, i.e. it would shift satiation-level curves $U_{AS}{}'$ and $U_{BS}{}'$ outside the box. So the key to the Hindu solution is that, due to the concavity of preferences, satiation points can be shifted outside the box as long as satiation-level indifference curves still cross within the box, including its edges; when a single crossing is left on the edge (which necessarily will be at one vertex of the box), all possible savings have been realized.

At this unique optimum $P$, god A receives all and only the available amount of $y$ and god B receives all and only the available amount of $x$, and both receive their satiation-level amounts. Interestingly, this solution implements a specialization pattern in which each god completely specializes in the offering for which his satiation level is lower than that of the other god, and therefore is "cheaper" to satisfy. As shown above, this is also one of the two specialization patterns possible in the specialized gods model, where, however, there is an alternative specialization pattern which does not allow any saving of resources—and one or the other pattern simply "happens" by assumption. By contrast, the Hindu solution does not assume specialization but concave preferences, and specialization is the ultimate result of concavity plus saving of excess offerings.

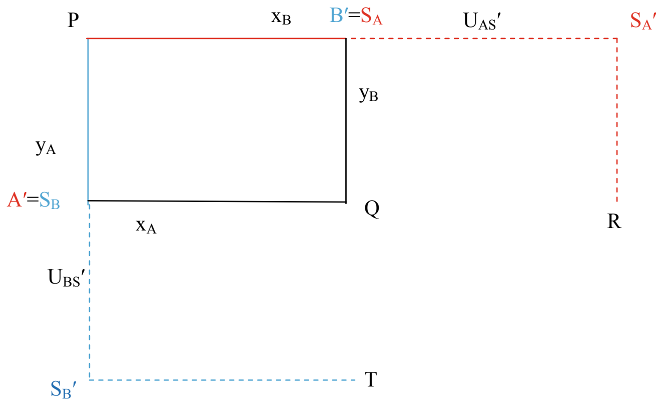


**Fig. 8.4** Sectarian Hinduism: saving religious resources

However, consider a variant of the model (not shown) in which satiation point bundles are such that one god desires lower levels than the other god in both offerings, i.e. $x_A^S > x_B^S$ and $y_A^S > y_B^S$. God B is here a junior partner of god A, with more modest claims on both offerings. Then god A's satiation bundle dictates the size of the box, his satiation point $S_A$ coincides with vertex $B$, and god B′ satiation curve $U_{BS}$ is entirely contained within god A's satiation curve $U_{AS}$, unlike in Fig. 8.3. The reader can easily check that now the box can be downsized in two ways to eliminate the excess offerings, either by squeezing it from the left or from below, yielding a unique optimum allocation in each case; at this optimum, however, in the former case god A completely specializes in offering *y* and god B in *x*, while in the latter case the opposite specialization pattern occurs. It is easy to understand the reason for the difference: in this variant, unlike the one depicted in Figs. 8.3 and 8.4, junior god B is cheaper to satisfy than god A on both counts, so as long as he is to be retained in the pantheon at a Pareto-optimal allocation, there is no special gain to be had from granting him specialization in one offering rather than the other.

Admittedly, at any of the unique optima just discussed, the supplicant is worshiping both gods, though only one god in each field. As it stands, the model cannot handle complete dedication to one god in both fields, but it suggests how concave preferences can overcome divine rivalry, which is the essential point at issue. This in turn allows not only the resolution of the supplicant's dilemma but also the saving of resources previously committed to religion, and thus the attainment of a superior religious outcome for the supplicant. In this framework, a way of rationalizing the worship of both gods may be to think of one of the gods as a junior partner in a sect devoted to the other god, with his own limited jurisdiction—for example, Shiva in a Vaishnava sect (as mentioned in Sect. 6.1.1). The last-discussed variant of the Hindu model seems particularly suitable for this interpretation, as it allows the "choice" between specializing junior partner B in one offering or the other, and thus lends itself to alternative applications—for example, the inverse positioning of Vishnu in a Shaiva sect. Moreover, as we have seen, despite the enormous distance in theological sophistication and in historical context, "implicit" Norse theology near the end of the pagan period was evolving toward an outcome not so different from the Hindu one, although embryonic and not supported by a sophisticated priestly science but carried by the simple people's worldview and embodied in their worship practices. The Vikings' perceived freedom to "choose" one deity as one's all-purpose friend and protector, while not denying respect to the other deities in a secondary way, ultimately boils down to overcoming the fear of divine jealousy and making the jurisdictional overlap harmless, thus escaping from the supplicant's dilemma. The concave preference model, especially in its junior partner variant, seems adequate to accommodate this too.

The Zoroastrian response to the jealous gods problem of traditional polytheism was the creation of a hierarchical pantheon. By the new doctrine, Ahura Mazda is supreme and has an all-encompassing jurisdiction, while all the other divine beings—the great Amesha Spentas and the other *yazatas*—were created and appointed by him to preside over a well-defined field, without encroachment upon one another's jurisdiction; that is, they are specialized deities. In terms of religious history, the old pantheon of *yazatas* became organized around the Amesha Spentas introduced by Zoroaster as helpers, cooperating with one another and with the supreme Lord for a common end and thus eliminating any rivalry (Boyce 1979, 41). As such, a *yazata*'s satiation level can be captured by a point on "his" side of the Edgeworth box, denoting the offering he cares for, while the other offering is considered as "neutral". By contrast, Ahura Mazda can still be depicted by strictly convex preferences culminating at an interior satiation point.

Figure 8.5 depicts the Zoroastrian theology, with the same box size and the same satiation point $S_A$ for god A—now the supreme god Ahura Mazda—as in Fig. 8.1. By contrast, god B—now a *yazata*—cares only about good $x$ and considers good $y$ as "neutral"; hence his indifference curves are straight vertical lines starting from each point on the upper side of the box. His satiation bundles are described by point $S_B$ and by the indifference line $U_{BS}$ starting from it; this line passes through the satiation point $S_B$ of Fig. 8.1, meaning that his satiation level of $x$, $x_B^S$, is unchanged. For clarity, an indifference curve $U_{A2}$ is drawn to capture A's utility level at B's satiation point. As a consequence of these changes, the contract curve that in Fig. 8.1 crossed the SE quadrant, $S_AP\ S_BQ$, is replaced by the thickened horizontal segment $PS_A$, which is the set of efficient allocations representing the tradeoff the supplicant now faces between satisfying the two divinities. Thus scarcity of resources is still there, implying that the supplicant cannot satiate both divinities, but now it involves only offering $x$.

In Fig. 8.5, good $y$ is in excess supply: starting from the allocations on $PS_A$ and adding any quantity of $y$ comprised in the quadrant $BS_BPS_A$ to a given quantity of $x$ leaves both divinities indifferent. This immediately suggests that the box is inefficiently oversized as it ties down excess resources to offering $y$. If resources committed to offering $y$ can be costlessly shifted to offering $x$, leaving the total expenditure unchanged, this can be improved upon. In Fig. 8.6, the box is shortened in height and lengthened in width with respect to Fig. 8.5 in such a way that the total expenditure of religious resources $R$ is unchanged (the horizontal segment $BB'$ in Fig. 8.6 is equal to the vertical segment $PS_B$ in Fig. 8.5), thus turning all the excess

**Fig. 8.5** Zoroaster's divine hierarchy: Ahura Mazda (**A**) and a *yazata* (**B**)

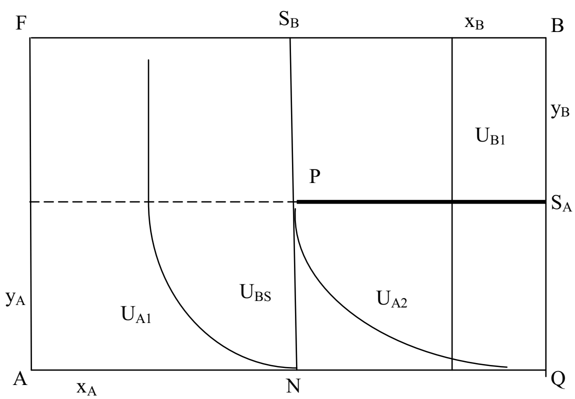

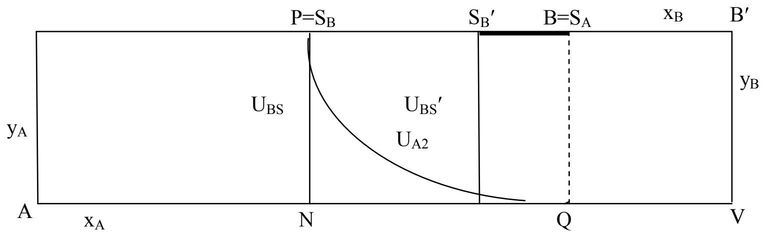


**Fig. 8.6** Efficient allocation of worship under divine hierarchy

amount of $y$ into additional $x$. A's position is unchanged, except that his satiation point $S_A$ is now found on the upper side of the box, coinciding with previous origin $B$. B's geometry is, however, changed because his point of origin is moved from the original $B$ to $B'$ and therefore his satiation point is moved from the previous $S_B$ (reproduced here from Fig. 8.5 for clarity) to $S_B{}'$. As a result, $S_B{}'$ still lies to the left of $S_A$ but is closer and scarcity has been mitigated, as has the supplicant's tradeoff, by the additional $x$ (the thickened segment $S_B{}'S_A$ is shorter than the segment $PS_A$ in Fig. 8.5). However, depending on how large was the horizontal distance between satiation points $S_A$ and $S_B$ under polytheism (in Fig. 8.1) relative to the size of the box, the reallocation from $y$ to $x$ under Zoroastrianism might well be sufficient not just to reduce but to completely eliminate scarcity and achieve satiation of both divinities, thus doing away with the tradeoff (the new $S_B{}'$ would then lie to the right of $S_A$). In any case, the general result is that this reallocation—made possible by the reduction of B from god with encompassing jurisdiction, overlapping with A's, to specialized, subordinate yazata—allows the saving of resources previously tied to pleasing everyone on everything and thus the attainment of a superior religious outcome for the supplicant, reducing—even when not eliminating—his dilemma.

This conditional result seems apposite. As we have seen, unlike traditional Iranian polytheism (and its Indo-European counterparts), Zoroastrianism is a strenuous faith, laden with moral obligations, purity laws, and ritual observances. As a consequence, the expenditure of resources that determine the size of the Edgeworth box, and which may be in scarce supply compared to the full demands of the faith, here must be understood as opportunity costs. These include not just the direct cost of the offerings (and the upkeep of the priests) but also the value of the time and effort that the supplicant is asked to devote, and of the consumption that he is asked to forgo, for the discharge of his individual duties—duties which are particularly testing and time-consuming in this religion. So the supplicant may not be able to fully live up to the demands and may again be forced to submit to a (reduced) choice dilemma.

Two notes may usefully conclude our discussion. The first regards the flexibility of use of religious resources. We have assumed throughout that total offerings devoted to each outcome can be made variable subject to an overall resource constraint $R = \overline{x} + \overline{y}$, but we have used this property only in the Zoroastrian case, because in the Hindu case the saving of resources involves no shift from $y$ to $x$ but only reduction of both (cf. Fig. 8.4). Suppose now that for whatever reason, $\overline{x}$ and $\overline{y}$ are fixed offerings in kind, so they cannot be converted from one use to the other—for example, offering $x$ is required to be personal attendance at public rituals while offering $y$ consists of spending time undergoing purification in isolation, with no substitution allowed. How would the Zoroastrian type of solution be affected? Clearly, there could be no restructuring of the box and no Fig. 8.6, hence no way to resolve or mitigate the supplicant's dilemma over offering $x$. However, waste could still be avoided simply by squeezing the box from above, in Fig. 8.5, until $S_B$ coincides with P and B with $S_A$, getting rid of the excess offering $y$. This useless amount of $y$ could not be turned into much-needed $x$ but it would be saved for non-religious uses, which still represents a Pareto-improvement.

Finally, we have assumed throughout gods with overlapping, but non-coincident, jurisdictions. What if those jurisdictions were perfectly coincident? The two gods would in effect be identical twins, with identical satiation points as seen from each god's point of origin ($x_A^S = x_B^S$ and $y_A^S = y_B^S$), so by our rules of construction of the Edgeworth boxes, discussed at the beginning of the preceding subsection, the satiation point of each god would coincide with the origin vertex for the other god. It will be no surprise, then, that the rivalry between such heavyweights will be especially hard to manage. The reader can easily see the effects of this identity assumption by looking at our figures. In Fig. 8.1, the classical rival gods, there would be no room for "neutrals", all indifference curves would be strictly convex throughout the box, and the contract curve would cross the whole box from vertex $A$ (=$S_B$) to vertex $B$ (=$S_A$). In the Hindu case, with concave L-shaped preferences (Fig. 8.3), god A's satiation curve $U_{AS}$ would coincide with the upper and right sides of the box $FBQ$, while god B's satiation curve $U_{BS}$ would coincide with its lower and left sides $FAQ$, yielding two optimal allocations where these two curves meet, i.e. at the vertices $F$ and $Q$. So the supplicant's dilemma would still go away, but no saving of resources would be possible because any squeezing of the box from any side would not only push the satiation points out of the box, which as we have seen is all right in the concave case, but would prevent the satiation curves from meeting, thus wiping out all the optimal allocations. Finally, following Zoroaster's reform of the polytheistic system (Fig. 8.5), the *yazata*'s satiation point $S_B$ would coincide with vertex F, so the contract curve would become the whole upper side of the box, going from $S_A$ to $S_B$. However, since there is no excess offering of $y$ to drop or shift, there would be no way to reduce the supplicant's dilemma or to reduce the expenditure of religious resources inherited from the polytheistic system. On reflection, this last finding is neither surprising nor disappointing: it only goes to confirm that the whole Zoroastrian reform makes sense only if god A was not equal to god B under paganism, but was already in some ways superior to his rival, and precisely for this reason he was then chosen as supreme god (as we have seen in Sect. 4.2.1).[5] More generally, the whole problem of identical gods may largely be an artificial construct: if two gods had been really perceived as

[5] In the same vein, when god B was diminished to a *yazata* role by Zoroaster, he was advisedly given jurisdiction over offering $x$, not $y$: $x$ is the offering for which his satiation claim under polytheism (in Fig. 8.1) was lower than for $y$—it was "cheaper" to make him lord of $x$. Had he been given jurisdiction over $y$, no restructuring of the box and saving of resources would have been feasible.

identical in every respect, people would have long since equated them as a single divinity and thus gotten rid of the imagined rivalry, as so often happened in the Greek and Roman world.

## 8.2 Rent-Seeking or Rent Dissipation: The Priestly Dilemma

The divine preference model of the previous section considered a direct transaction between a supplicant and the gods and made no mention of priests; we started with the benchmark case of rival gods and then proceeded to consider different ways to escape the supplicant's dilemma.[6] In most of the polytheistic systems under study, however, priests mediate these transactions, so in the same theological setting where people are transactional in their relationship with gods and gods have overlapping jurisdictions, we now introduce the priests. Multiple gods may—or indeed, should—be propitiated for the same reason, but propitiating gods is a costly business in a world with scarce resources. Priests receive rents for mediating between gods and people. Therefore, propitiating one god reduces resources (and therefore rents to priests) available to propitiate another. We model a simultaneous game between two priests in this polytheistic setting, focusing on the contrast between Greco-Roman and Hindu priests; the other polytheistic systems will be briefly taken up toward the end.

In this model, each priest can choose between a "systematic" religion (S) and an "accommodative" religion (A), defined as follows. Systematic religions focus on ritualistic consistency and purity rather than any particular god. Priests in a systematic religion are effective at mediating transactions between gods and people precisely because *all* priests follow the same or very similar rituals. If they do not, then bad things can happen, or at the very least good things may not happen when they are needed. Even though gods have overlapping jurisdictions, the focus on consistent rituals implies that ritual propriety is critical rather than which god is being propitiated; therefore, all available resources are used to satisfy ritualistic consistency—the identity of the targeted god is not consequential. Hence, ritualistic unity *makes* cultic transactions non-rival: these gods are not jealous of each other. Thus, systematic religions capture an essential element of classical Hinduism, non-rival gods, complementing the model of divine preferences outlined above. By contrast, accommodative religions do not require a unified ritual set but are focused on the specific god. Consequently, propitiating one god through one ritual does not satisfy another who demands his/her own ritual; hence, given overlapping jurisdictions, all gods with similar jurisdictions must be propitiated. With scarce resources, gods with overlapping jurisdictions are jealous of each other because they all demand a share of the limited earthly resource pie. As we have seen, this is a crucial feature of Greco-Roman religion.

[6] This section is based on the model developed in Basuchoudhary, Ferrero and Lubin (2020).

Each priest makes his choice at the same time as the other, capturing the idea that at any given point in time, belief systems are private information. This "choice" could be thought of as implying meta-preferences over beliefs, but it can be rationalized in an even simpler way as a choice between giving priority to the service of a given god rather than another or to the ritual forms and practices of the religious service—something which does not actually involve faith. Be that as it may, the choice of S or A becomes a strategy for each priest, so there are four possible outcomes in this model: S-type priests can coordinate with each other, and A-type priests can do the same; however, it is also possible for one priest to choose S while the other chooses A, and vice versa. These interactions have payoffs to each player, and the game is symmetric in the strategy space.

If a systematic priest can coordinate with another systematic priest who believes in the same set of rules, then they can enforce their rituals for all transactions. This is what Brahmanical priests did, enforcing the Shrauta rituals laid out in the Vedas and developed in the late-Vedic literature. These rituals might invoke different gods for different or similar reasons, but they could only be invoked by Brahmins using similar, codified rites. In the age of sectarian Hinduism, the gods again had overlapping jurisdictions; for example, Shaivites might propitiate Shiva to receive the same benefit that Vaishnavites might seek from Vishnu. In contrast to the earlier Vedic age, Shaiva and Vaishnava priests no longer performed the same rites but there remained a structural similarity of rituals across traditions because they stemmed from the same Shrauta foundation and were tailored to the sectarian loyalties. Therefore, in this setting, which god is propitiated is secondary as long as uniform rituals are followed, and devotees of a particular god do not have to propitiate multiple gods to get a benefit. Scarce earthly resources do not have to be divided across many gods, and therefore across many priests and the concomitant transactions costs, to secure a specific benefit. The highest possible *total* rents/payoffs, therefore, occur with the {S, S} strategy profile.

Accommodative priests, on the other hand, do not need other priests to coordinate on rituals, as ritual unity and consistency is the essence of the S religion but is irrelevant to the A religion. Therefore, the very act of challenge to a unified ritualistic system lays out the possibility that those systematic rituals are ineffective. S priests cannot really coexist with A priests then because the latter's choice is a blow to S priests, implying that an A priest gains market share, and rents, at the S priest's expense. Thus, in the {S, A} or {A, S} strategy profile, the A priest receives a larger share of rents than the S priest.

Of course, accommodative priests, by definition, can coexist with other accommodative priests. This is the {A, A} strategy profile. Given overlapping jurisdictions among the gods they serve, scarce resources imply that less is available to propitiate individual gods and, therefore, to reward individual priests in this strategy profile. Warriors hoping for victory will ask priests to propitiate both Ares and Athena for fear of angering one or the other if s/he was left out. However, in a world of limited resources, the priests of Ares and Athena have a lower share of available resources than when all priests follow ritualistic purity, and gods do not compete with each other.

**Table 8.1** The priestly dilemma

| | S | A |
|---|---|---|
| S | 5, 5 | 0, 6 |
| A | 6, 0 | 2, 2 |

Our model, therefore, follows the structure of a simultaneous coordination game. We have structured the payoffs, in the light of the above description, like a Prisoner's Dilemma. The model does not have to be a Prisoner's Dilemma, as the religious structure described above may well be an Assurance game. However, keeping the payoff to deviating from S larger than the payoff to coordinating on S (6 > 5) allows us to discuss the possible mechanisms through which the Brahmanical system was resilient while the Greek and Roman system was not.[7] Table 8.1 represents this model. Obviously, in this setting, the Nash equilibrium is to accommodate—the {A, A} profile. Nevertheless, a simple application of the Folk Theorem implies that if the players are sufficiently patient, then the {S, S} profile is a subgame-perfect equilibrium. That is, if the game is repeated with some certainty $\delta$, then {S, S} becomes an equilibrium outcome for a large enough $\delta$.

Vedic rituals remained unwritten for perhaps a millennium; by definition, then, only initiates within hearing distance of a teacher could learn the rituals, and this limited the number of priests who could perform them. Moreover, there was an extreme emphasis on fidelity of transmission: rituals passed verbally from teacher to pupil with a great deal of certainty, and this continued even after the Vedas and later liturgical texts were written down. At the same time, the emphasis on *grihastha* (the disciplined householder) ensured these rituals were transmitted within families. Together, the reliance on auditory fidelity and a familial priesthood also created barriers to entry, which helped to keep rents high in the {S, S} equilibrium. As a result of these features, $\delta$ was high in the Brahmanical setting. It insulated the Hindu religion from accommodating pluralism by ensuring {S, S} as SPE.[8]

The priestly dilemma outcome was different in the Greco-Roman world. Priests had no incentive to teach rituals to anybody; priestly functions were official, and

[7] What would happen if the payoff to deviating from S were smaller than the payoff to coordinating? The game would be an Assurance game with a payoff-dominant equilibrium and a risk-dominant equilibrium. Here too, the relatively closed nature of the Brahmanical priesthood would create more particularized trust among them while ritual purity might generate generalized trust, leading to the payoff-dominant equilibrium. By contrast, the very accommodation among Greco-Roman priests and the lack of any focal ritual would potentially push their system toward the risk-dominant equilibrium. However, we chose the Prisoner's Dilemma version because we wanted to model the extreme case of large benefits to deviation. Thus, the model is hardwired with the sense that accommodating religions can bring large benefits to practitioners at the expense of systematic religions. If the systematic religion can survive this strong incentive to deviate merely strengthens our case.

[8] This inference is consistent with the idea that finitely lived agents with overlapping generations can achieve mutually beneficial equilibria if their life spans and overlapping generations are long enough (Kandori 1992). While actual life spans for Brahmins may be hard to know, the long apprenticeship within a familial priesthood common in the period we consider certainly suggests a more than usual "overlapping generation"—something completely missing in the Greco-Roman framework.

offices could not be passed down from father to son (except in some ancient Greek priesthoods controlled by gentile families). The theology of this world encouraged the dissipation of scarce resources in the propitiation of many gods to achieve the same purpose; so it is not surprising that, as we have seen (Sects. 2.1.3 and 2.2.3), the monetary incentive to even be a priest was low or negative (a cost, not a reward), and the value of priesthood was more a matter of signaling civic virtue and elite status than garnering rents. Thus, even if a priest was serious about his god, death would ensure a finite time horizon for the priest while the likelihood of repetition of the game (δ) was close to zero. We know that in this case {A, A} is the Nash equilibrium.

As we have seen, the rise of Shramanic movements made a dent in the Brahmanical priesthood, who lost significant followings and rents to them. However, the very intergenerational ritualistic continuity within the Brahmanical priesthood, coupled with the higher rents associated with a non-rival polytheism, inoculated Brahmanism from the competition by increasing certainty about the future of the repeated priestly dilemma game and therefore the expected payoff to maintaining the priesthood relative to any alternative. The ascetic sects, scorning the rituals themselves and lacking the family network, could not match this forward-looking attitude. So the Brahmins did maintain the coordination on the systematic strategy, holding Jainism at bay and outlasting Buddhism in India while stemming the tide of Abrahamic evangelism when it came.

By contrast, the very nature of pluralism in the Greco-Roman world worked in the opposite direction. The ever-expanding accommodation of rival gods and the lack of a professional, hereditary priestly class preserving a unified ritualistic system ensured an uncertain future for a priest's job and thus tended to slacken his commitment to the group. It is hard to care about a god whose priesthood confers little benefit on the living priest, the priesthood itself, and the children of the priest. If those priests saw people drifting away to new gods or new religions, they might have complained about the loss of ancient virtues and the decay of the world but would not have taken action. Nevertheless the system was stable for many centuries because it was never really challenged. When, however, an alternative pattern of civic virtue was forcefully promoted by Emperor Constantine and his successors, there was little incentive for the elite to keep signaling civic virtue by holding on to a priesthood of Jupiter. And they did not, making a relatively smooth transfer from a priest of Jupiter to a bishop of Christ when the state required it.

In addition to modeling the formation of a priests' monopoly as the outcome of an indefinitely repeated Prisoner's Dilemma game among them, another, complementary approach is to see this monopoly as an established incumbent subject to a threat of competitive entry and engaged in an entry deterrence game with the potential entrant (Ferrero 2014b, c). Two exclusive religions employ missionary effort to maximize the number of members (or converts) from a fixed population, like two boats that go out fishing from a common pool; membership maximization can be seen either as an effort to acquire new members or to retain existing members who might otherwise drift away to the competition. The incumbent enjoys a first-mover advantage: it can credibly precommit to a given effort level if entry occurs (like a preemptive capacity expansion), which will decrease its marginal cost of effort; this precommitment can be such as to either deter entry, in which case the incumbent remains a monopolist, or accommodate the entrant in a Stackelberg equilibrium in which the incumbent is the leader. By making a prior decision on its effort commitment, the incumbent can in effect choose whether to let the entrant in or keep it out, whichever solution brings the incumbent higher benefits. The deterrence option turns out to be superior to the accommodation option if the entrant's entry cost is sufficiently high, and this cost can be further manipulated by the incumbent (like the building of an entry barrier) to its own advantage. Even if deterrence is optimal, however, the threat of entry distorts the incumbent's level of effort upward, and hence its benefits downward, relative to the outcome it would achieve if its monopoly went unchallenged. So in this equilibrium, entry is not observed but overcommitment of effort is. In the articles cited above, this entry deterrence model was applied, respectively, to Pauline Christianity' response to the threat of Jewish entry into the Gentile mission market and to the Catholic Church's response to the Protestant Reformation (the so-called Counter-Reformation). However, it could as well be applied to the Brahmins' response to the entry threat by the Shramanic sects, when the Brahmins engaged in an extraordinary effort to re-invent themselves and multiply and differentiate their services; and, conceivably, it could also be used as a model of a non-event, that is, a model of what could have happened if the Greco-Roman priests had managed to develop a corporate, unified response when faced with the Christian entry threat.

Our model of a priestly game is nicely complemented, and supported, by Tridimas' (2021) study of Greek religion. He asks if a priestly interest group, hence a monopolist supplier of religion, will come into being starting from a collection of single unorganized priests, and models the problem as a utility-maximizing decision by a priest about whether or not to join the group. The model shows that the interest group will fail to form if the cost of joining is greater than the benefit from the very start, and this will be the more likely, the higher is the number of gods worshiped and the higher the political power of the citizens. Focusing on the Greek case, the first factor increases the heterogeneity of the priestly class because Greek priests were specialized by god and temple, thus increasing the individual cost of joining, while it decreases the rents potentially available to the group through dispersion of resources, thus reducing the individual benefit. The second factor reduces the probability of successful rent extraction by the interest group because the Greek democratic city-states, unlike the kings of old, had no use for divine legitimation of their government. This is another way to reach the same result we found with the priestly dilemma game.

Turning to the other religions, we know next to nothing about the early centuries of Zoroastrianism and the role of the priests in it; we only know that at the end of the process, when its outcome is recorded in the Avesta, the role and prominence of the priesthood is greatly enhanced. We know something about the starting point though: the organization, recruitment, and functions of the Iranian priests were very similar to those of the Vedic priests; in contrast with the Indo-Aryans, they migrated to territories populated by kindred people and with broadly similar climate and resources, which lessened the problem of assimilating foreign local deities; and most importantly, as far as we know, they were never the target of challenges by ascetic sects or other outsiders as the Brahmins were. There surely was conflict with the original "accommodationists" among their ranks, i.e. the priests devoted to the banned gods (Indra and the other *daevas*) who must have been popular among Iranian peoples, especially the warrior groups. But then they enjoyed a unique bonus unavailable to their Indian colleagues: a rupture in historical continuity set in a precise, if unknown, moment in time when a prophet initiated the building of the new religion from the bricks of the old one. All these factors together would have made the job of (S, S) coordination a relatively simpler matter for the Iranians than for their Indian cousins, so if the priestly dilemma model is adequate to capture the Indians, it will be all the more fit to the Iranians.

On an a priori basis, one would think that the model should also fit the Celts, whose druids in Gaul were reported by the classical writers, particularly by Caesar, as having an organization strongly reminiscent of the Brahmins'—indeed, one even going beyond that in the sense that they met annually in a general council to decide matters and were presided over by a chief druid, something that the Indian priests never had. Unfortunately we will never know because we have no reliable information after Roman times; in particular, we know nothing about the Irish priests, who appear in the extant Christian sources only as sorcerers trafficking with demons that St Patrick and his successors readily defeated. So there is no way of knowing how, if at all, they confronted the Christian onslaught and why, unlike their Indian colleagues, they went down so thoroughly as to leave no trace. Indeed, we do not even know if their theology evolved in any way over time to address the supplicant's dilemma problem, even before the rise of Christianity. So the Celtic case must remain an open question.

By contrast, the Germanic and Scandinavian case most definitely lies outside the purview of our priestly model. As we have seen (Sect. 3.2.2), even toward the end of the pagan period, not to say earlier, the Scandinavian priesthood was not just fragmented, uncoordinated, and subordinated to politics like the Greek and Roman ones: it existed only scantily and embryonically, most priestly functions being fulfilled by chiefs and kings as a side occupation among their secular duties. So, even aside from the emerging non-rivalry of their theology, it is no surprise that Christianization of these peoples went through basically unopposed.

## 8.3 Institutional Equilibria: A Game Between Gods and Priests

The idea of institutional equilibria (discussed in Sect. 7.6 above) can be made more rigorous in the framework of a coordination game—a type of game which has multiple equilibria and where the occurrence of one equilibrium rather than another depends on the players' ability to coordinate their moves. The players here are the priests and the gods. This may seem odd as gods do not "play" a game—they are what they are, their nature and personalities being inscribed in the cult and mythology of the various religions surveyed in this book. Nevertheless, even though the gods' chosen "strategies" at any given time and place are the product of prior evolution, it may be instructive to think of the priests—these, more properly players—as playing their strategies against gods conceived in different ways in the different theologies and see what the outcome is. On reflection, the priests too are what they are depending on the religious organization in the different societies, but in the previous section we treated them as making strategic choices just as well. Here we consider not individual priests but the priestly class as a group, and similarly, not individual gods but the gods of a given theology as a group. Thus, our players are groups.

The priests' strategies are Corporate and Atomized, i.e. behaving as a corporate profession or as atomized practitioners of religious service. We change the terminology from the previous section because here we do not have a game among priests but one between priests and gods, so what counts is the institutional organization of the priesthood. Corporate priesthood was found with the Hindus (H), the Zoroastrians (Z), and the Celts (C), whereas atomized priesthood was found with the Greeks (G), the Romans (R), and the Vikings (V). The gods' strategies revolve around the most critical feature emphasized throughout our study: they can be Jealous or Non-jealous. As we know, gods are jealous with the Greeks, the Romans, and the Celts, while they are non-jealous with the Hindus, the Zoroastrians, and the Vikings. These strategies yield payoffs for the players. The payoffs of the two players are, strictly speaking, incommensurable: it is a rent for the priests and a level of utility for the gods; remember, however, that behind the gods' perceived satisfaction lies the very concrete satisfaction of the supplicant in his quest for benefits and protection from them. To avoid giving the impression that the two players are sharing in the "same" pie, we made the model non-symmetric, i.e. the players are not interchangeable.

In our model, there are two equilibria, one of which offers both players a higher payoff than the other; so if faced with a choice among equilibria, both players would agree on it. Achieving such a payoff-dominant equilibrium, however, requires coordination, but each player chooses his strategy independently and simultaneously with the other. If each player is uncertain about the choice of the other player and does not trust his willingness to cooperate, they may end up at the alternative, Pareto-inferior equilibrium because the payoff that each player achieves there does not depend on coordinating with the other player. This type of coordination game is called an Assurance game, exemplified by the classic Stag Hunt game.

Table 8.2 represents this model. By the usual convention, the pair of numbers in each cell gives the payoff of the Row player (the priests) and the payoff of the Column player (the gods), in this order; each cell also indicates the religion(s) where

**Table 8.2** The gods and priests game

| Priests | Gods | |
|---|---|---|
| | Non-jealous | Jealous |
| Corporate | HZ 6,5 | C 2,4 |
| Atomized | V 4,1 | GR 4,3 |

the payoffs occur. There are two Nash equilibria in pure strategies, (Corporate, Non-jealous) and (Atomized, Jealous), of which the former payoff-dominates the latter as both its payoffs are higher. Note that the payoff-dominated equilibrium (Atomized, Jealous) is "safer" in the sense that each player can achieve an equal or higher payoff if the other player deviates from that equilibrium (if Row plays Atomized while Column plays Non-jealous, he still gets 4; if Column plays Jealous while Row plays Corporate, he gets $4 > 3$); this is not true for the (Corporate, Non-jealous) equilibrium, which requires coordination between the players to be implemented. Additionally, to further emphasize this aspect, we gave the (Atomized, Jealous) equilibrium the property of risk dominance, a refinement of the Nash equilibrium concept introduced by Harsanyi and Selten (1988): this equilibrium is less risky, hence the more uncertain players are about the actions of the other player(s), the more likely they will choose the strategy corresponding to it.[9]

The relationships between the numbers in Table 8.2 are meant to capture the differences among the religions of interest, as discussed at length in this book—keeping in mind that only relative values (greater or lesser than or equal to) matter, not absolute values per se. If priests are a corporate profession (upper row), their rents are higher when they exercise their monopoly in the service of non-jealous gods (6), like the Hindus and the Zoroastrians, than when they have to appease a range of jealous gods (2), like the druids; on the other hand, if priests are atomized (lower row), their rents are about the same (4) when they are one with the social elite honoring non-jealous gods, like the Vikings, or when they are part-time specialized officials servicing jealous gods, like the Greeks and the Romans. From another angle, if gods are non-jealous (left column), priests fare better when they are corporate (6) than when they are atomized (4) because the corporate organization makes them central in the performance of cult; on the contrary, if gods are jealous (right column), priests fare better when they are atomized (4) than when they are corporate (2) because the druids of the latter structure, unlike the Greco-Roman priests of the former, were required for the performance of all sacrifices and hence, presumably, under a much greater stress to please all the gods with their idiosyncrasies.

Turning to the gods, non-jealous gods (left column) attain a far higher utility when priests are corporate (5) than when priests are atomized (1) because in the former case, following the Hindu and Zoroastrian reforms, the supplicant' dilemma is eliminated or reduced whereas in the latter case the Viking public cult is kept to modest proportions; on the other hand, jealous gods (right column) presumably received a somewhat more systematic cult when managed by corporate priests like the druids (4) than when serviced by the fragmented Greek and Roman priesthood

[9] Technically, a strategy pair risk-dominates another if the product of the two players' losses from deviation is higher for the former than for the latter. If Row plays Corporate and Column deviates to Jealous, Rows gets 2; if instead Row himself deviates to Atomized, he loses 2 if Column plays Non-jealous but avoids losing 2 if he plays Jealous. At the same time, if Column plays Non-jealous and Row deviates to Atomized, Column gets 1; if instead Column himself deviates to Jealous, he loses 1 if Row plays Corporate but avoids losing 2 if he plays Atomized. So the product of the losses from deviation is $(2–4)\ (1–3) = 4$ for (Atomized, Jealous) while it is $(4–6)\ (4–5) = 2$ for (Corporate, Non-jealous): the former strictly risk-dominates the latter.

(3). From another angle, if priests are corporate (upper row), gods achieve a greater satisfaction when they are non-jealous and the supplicant's dilemma is resolved (5) than when they are jealous and the dilemma is in place (4); on the contrary, if priests are atomized (lower row), gods are better serviced when they are jealous (3) than when they are non-jealous (1) because in the former case the public cult is targeted to controlling the gods' touchiness whereas in the latter case the anxiety about the gods' possible neglect is allayed so that the pressure to worship is reduced.

If, then, the relationships among the values in Table 8.2 make sense as stylized descriptions of the fundamental character of theology and priesthood in our religions, our modeling approach gives a crisp meaning to the idea of alternative institutional equilibria: they are the two Nash equilibria in an Assurance game of gods and priests, and they are "institutional" because they combine divine preferences with priestly institutions. Precisely because they were equilibria, i.e. configurations toward which the religious system tends to gravitate, these alternative arrangements were stable and long lasting: in both of them, the priesthood structure was congruent with the theology. In the (Corporate, Non-jealous) equilibrium, the Hindu and Zoroastrian priests managed to resolve the quandary that beset Indo-European polytheism, make the gods non-jealous, and thereby achieve an efficient allocation of religious resources and maximize the rents for themselves – which is why we called it a payoff-dominant equilibrium. In the (Atomized, Jealous) equilibrium, the Greco-Roman priests submitted to the pervasive and growing problem of overlapping jurisdictions among jealous gods by adjusting to it and taking a low profile, making religious service a part-time, specialized business, and forgoing any prospect of exercising a monopoly power and claiming the associated rents. The stability of such an arrangement, despite its ubiquitous and growing inefficiency, was due to inertia, or the lack of any viable alternative and any serious challenge before the Christian onslaught in the Roman Empire. There is a sense in which, absent a decisive effort at reform on the part of the priestly class, this Pareto-inferior equilibrium was the natural outcome of the evolution of Indo-European polytheism as we have described it: an outcome that requires not action but inaction, and which therefore is safer—which is why we called it a risk-dominant equilibrium.

Conversely, the model makes clear why the Viking and the Celtic arrangements failed to take hold and proved brittle and transient: they were not equilibria of a coordination game, because in both, even if in opposite ways, the theology and the priestly institutions were mismatched. The Vedic Brahmins, who were apparently so similar to the early druids, may have been in a similar predicament, described by the cell (Corporate, Jealous) which is out of equilibrium; but then, under the pressure of competition from the ascetic groups, they found the strength and the inventiveness to turn the theology around and land on the (Corporate, Non-jealous) equilibrium. Similarly, the earliest Greeks too—the Mycenaeans, and perhaps their successors in the Dark Age from which no information survives—may have had an aristocratic, professional, hereditary priestly class linked with the royal palace, which could perhaps be housed in the same cell as the druids and the Vedic priests since they pantheon was already oversized and conflict-ridden. But then, under the pressure of the new democracy of the polis and the obvious difficulty of handling an ever-growing pantheon, they gave up any residual of corporate priesthood, took the line of least resistance, and landed on the safer (Atomized, Jealous) equilibrium.

## 8.4 The Missionary Expansion: A Discriminating Cooperative Model

For ease of exposition, we first set up and solve the mission problem with reference to the Zoroastrian mission in the Iranian setting.[10] We then show that, by suitable reinterpretation, the model can also be used to rationalize the Brahmanical expansion in the classical Hindu period.

Traditional polytheistic religions are not missionary enterprises and one does not "convert" to them (except perhaps spouses and slaves), so we can think of traditional Iranian religion as a religious community that from time immemorial had structured itself in such a way as to provide the maximum net benefits to its members; alternatively, one that had acquired a level of membership that was efficiently maximizing net benefits, or welfare, per capita. The switch to Zoroastrian monotheism and the start of a mission to convert other Iranian peoples inevitably involved, on the one hand, a fall in per capita benefits because the community was diluted and the priests distracted toward missionary work, and on the other hand an increase in the marginal and average cost borne by members, as new members were naturally more and more difficult to convert and retain as the expansion proceeded. Nevertheless, the switch to mission—the road that ultimately led to an empire-wide religion—could be made acceptable to the original community and still remain viable if the *total* net benefits generated by the *new* members were sufficient to both compensate the old members *and* leave a residual—a rent—to support the expanded priestly class that the missionary spread of Zoroastrianism, as we have seen, entailed. Thus described, the move involved the equivalent of a kind of wage discrimination—unequal post-transfer benefits for old and new members.

If we think of religious consumption as the output of a household production process that employs only the members' "labor" as an input, the traditional community equilibrium described above can be modeled as the solution to the problem of a producer cooperative that chooses its membership level to maximize net benefits per member—measured as the difference between gross benefits and cost of participation. Starting from here, expansion yields net benefits from new members; these can be partly siphoned off to compensate the original members and partly used to provide a rent to the new priests, while still leaving a nonnegative residual net benefit to the new members themselves. Hence, the new equilibrium level of total membership is constrained by the condition that the total net benefits generated by the new members be strictly greater than the total losses of the old members. However, an even better solution is one that turns this constraint into an objective, i.e. one which maximizes the difference between net benefits from new members and losses of old members.

---

[10] This section is based on the model developed in Ferrero (2021).

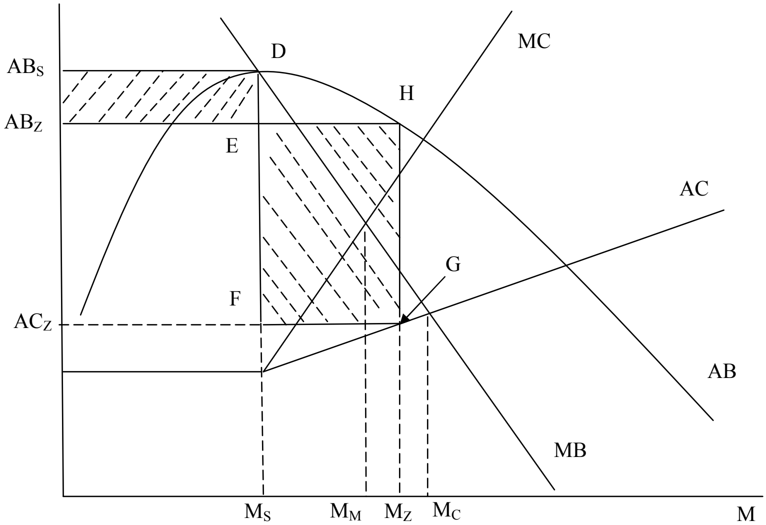


**Fig. 8.7** The Zoroastrian mission and the priests' rents

Figure 8.7 depicts the Zoroastrian mission problem and its solutions. The above assumptions imply that average benefits are first increasing and then decreasing with membership, and that the optimal membership of the traditional community occurs where average benefits (*AB*) peak, i.e. where they equal marginal benefits (*MB*) from additional members, determining membership $M_S$ and benefits per member $AB_S$. The subscript S denotes the equilibrium of a traditional community that functions as a closed sect which maximizes its members' average benefits; this solution is analytically identical to the classic solution of a dividend-maximizing producer cooperative.[11] Missionary expansion, following the Zoroastrian reform, starts from here. As posited above, unlike with the old members (who have long since adjusted to sect life and can be taken to be homogeneous), the average cost (*AC*) and marginal cost (*MC*) of joining are assumed increasing with new members because, at least beyond a certain point, conversion involves people who are more removed from the original group and whose opportunity cost is therefore higher.[12] Seen from another angle, the Zoroastrian community enjoys potential market power as it is the only

[11] The theory of the producer cooperative or labor-managed firm is old but seems now out of fashion. For a good introduction to the model see the survey by Bonin and Putterman (1987) and the literature cited therein. A full analytical treatment is in Ireland and Law (1982). Nowhere in this literature, however, is our special constrained optimization problem addressed.

[12] Although we have next to no information on the prehistoric spread of the religion, it may well be that for an initial range of expansion the average cost of membership would have decreased, for example because of a fixed cost or of network effects. The Appendix to Ferrero (2021) shows that, under some conditions, the Zoroastrian solution can survive this extension. Unlike in the increasing cost case, however, it may (though need not) imply a level of membership lower than in the "monopsony" solution because the rapid expansion of *M* triggered by the decreasing *AC* also involves a rapid fall of *AB* and hence a large loss for the old members.

provider of that brand of religion, and so it faces a supply curve of members which will sooner or later slope upward; symmetrically, the marginal benefit curve *MB* can be thought of as its demand curve for members. Even through its expansion process, however, it remains a cooperative organization bound by the constraint to protect the welfare of the original group—in effect, a discriminating cooperative that redistributes benefits from new to old members. To proceed further, we need some simple math.

With little loss of generality, we use quadratic average and marginal benefit functions and linear average and marginal cost functions to derive easily comparable closed-form results. Let $TB = \alpha M^2 - \beta M^3$ (with $\alpha, \beta > 0$) be the religious community's total benefits as a function of membership *M*. This yields average benefits per member:

$$AB = \alpha M - \beta M^2 \tag{8.3}$$

and marginal benefits:

$$MB = 2\alpha M - 3\beta M^2 \tag{8.4}$$

Function (8.3) has an interior maximum, which is found by maximizing it with respect to *M* and coincides with the value of *M* that equates (8.3) and (8.4) (see Fig. 8.7). This yields the level of membership $M_S$:

$$M_S = \alpha / 2\beta \tag{8.5}$$

This is the standard solution of a producer cooperative that determines its membership so as to maximize benefits (or income) per member, and will be the solution that describes our traditional community equilibrium. Note that this solution is not responsive to the availability of outsiders who might be willing to join the community to partake in the benefits (for a producer cooperative, the labor supply), for, on the traditional assumption that all members receive the same benefits, their admission would lower the existing members' average benefits.

When the community starts on its missionary expansion, the outsiders' average opportunity cost of joining (in production, the labor supply price or wage) becomes relevant. This average cost is constant, and lower than average benefits, for the old members up to $M_S$ and then increases with every new recruit as $AC = \gamma M$, $\gamma > 0$, for simplicity; the corresponding marginal cost is $MC = 2\gamma M$. Hence both the *AC* and the *MC* curves cross the *AB* curve in its decreasing region, which implies that, at the starting equilibrium $M_S$, there are people whose cost of joining is lower than the current average benefit level, so they are willing to join.

As benchmarks, it is useful to compute two standard solutions. The first is the solution that maximizes *total* benefits net of the cost (in production, total profits) without exploiting the community's market power but taking average cost as if it were a market parameter (like a market wage)—in effect, the "competitive" solution *C*. This solution is found by equating *AC* to marginal benefits *MB* (Eq. 8.4) and yields membership $M_C$:

$$M_C = (2\alpha - \gamma)/3\beta \tag{8.6}$$

The second benchmark is the solution that maximizes total net benefits exploiting the community's market power vis-à-vis potential new members—in effect, the "monopsony" solution *M*. This is found by equating *MB* (Eq. 8.4) to marginal cost *MC* and yields membership $M_M$:

$$M_M = 2(\alpha - \gamma)/3\beta \tag{8.7}$$

Obviously, as can be easily checked, $M_C > M_M$.

As explained above, the condition for the missionary expansion to be both viable and acceptable to the traditional community is that the *total net* benefits brought in by the *new* members be strictly greater than the total losses incurred by the *old* members from the lowering of their traditional benefits: this is the *compensation constraint*. If the former is greater than the latter, it allows for full compensation of the old members while still leaving a positive residual to both finance the missionary expansion and provide nonnegative net after-tax benefits to the new members. At any membership level $M > M_S$, the difference between these two measures is:

$$(AB - AC)(M - M_S) - (AB_S - AB)M_S \tag{8.8}$$

where $AB_S$ is found by substituting (8.5) into (8.3).

Calculation shows that (8.8) is greater than zero at both $M_C$ and $M_M$, satisfying the compensation constraint. However, the community can do better than either and turn the compensation constraint into the *objective* of a maximization problem. Since inequality of benefits between old and new members is necessary for the former to agree on the mission, if we think of the community as seeking the most profitable expansion over and above full compensation of the old members, this is tantamount to maximizing (8.8) with respect to *M*. Using the above expressions to substitute into (8.8), the FOC for a maximum is:

$$2(\alpha - \gamma)M - 3\beta M^2 + (\alpha\gamma)/(2\beta) = 0 \tag{8.9}$$

which can be rewritten as:

$$2\alpha M - 3\beta M^2 - 2\gamma M = -(\alpha\gamma)/(2\beta) = -\gamma M_S < 0 \tag{8.9$'$}$$

or:

$$2\alpha M - 3\beta M^2 - \gamma M = \gamma M - (\alpha\gamma)/(2\beta) = \gamma(M - M_S) > 0 \tag{8.9$''$}$$

The positive root of (8.9) yields the Zoroastrian solution $M_Z$:

$$M_Z = \left[2(\alpha - \gamma) + \sqrt{4(\alpha - \gamma)^2 + 6\alpha\gamma}\right]/6\beta \tag{8.10}$$

Direct comparison of (8.5), (8.6), (8.7), and (8.10) shows that $M_S < M_M < M_Z < M_C$. Expressions (8.9′) and (8.9″) provide analytical proofs of these results. The LHS of (8.9′) is $MB - MC$, which is equal to zero at $M_M$ but negative here, proving that $M_Z > M_M$. The LHS of (8.9″) is $MB - AC$, which is equal to zero at $M_C$ but positive here, proving that $M_Z < M_C$. Thus the Zoroastrian solution turns out to lie somewhere in between the competitive and the monopsony solutions, as shown in Fig. 8.7.

As can be seen in the figure, the move from $M_S$ to $M_Z$ (or to any other level of $M$ greater than $M_S$) would not be acceptable to the old members without redistribution and discrimination because average benefits fall from $AB_S$ to $AB_Z$. The outcome at this equilibrium is shown by the two shaded rectangles: the area EFGH measures the net benefits from the new members while the area $\mathrm{AB_S}$ DE $\mathrm{AB_Z}$ measures the total losses of the old members; the difference between these two areas, though positive also at other membership levels such as $M_M$ and $M_C$, reaches a maximum at $M_Z$. This confirms that the switch to missionary monotheism can be Pareto-improving and therefore unanimously accepted.

We have hardly any direct observation of the missionary, pre-state period of Zoroastrianism, so one wonders what the compensation to the old members may have been then. Enhanced reputation and influence, which facilitated profitable trade connections in the newly converted territories, are a fair guess (cf. Boyce 1982, 7–9, for the spread of Zoroastrianism in western Iran). In the longer run, however, there was one great new benefit: the fire temples, which began under the Achaemenians and spread all over the empire, including the northeastern region of the Iranian plateau which was the homeland of the original Zoroastrian community. These "old members" surely drew benefits from such institutionalization of the cult—witness the fact that the fire temples became a fixture of Zoroastrian communities the world over, down to the tiny groups surviving today.

Turning to India, we have seen that the Brahmins' obsessive concentration on their own uniqueness and separate identity as a class allowed them to overcome the dark period of the ascetics' rise and the concomitant weakening of the traditional royal demand for their *shrauta* services. They managed to successfully promote themselves as ascetics of a new kind and at the same time superior providers of non-religious services to the royalty and other elite. This happened through the Brahmins' migration to new lands in South India and outside the subcontinent, often at the behest of local rulers; hence, it must have involved a growth in the numbers of "active" Brahmins, i.e. those who not only qualified as such by their training and lifestyle but also found employment as professional providers of the above services. We can then call this movement a missionary movement, on the understanding, however, that their mission was not about converting anybody to new gods or theological beliefs—although they did carry the new Hindu sects with them and helped to assimilate local, foreign deities into those sects—but about "converting" people to Brahmanism itself, i.e. the ideology of Brahmanical supremacy. At the same time, this renewal and elevation was conditional on their ability to maintain a tight, closely guarded monopoly of their services and prevent uncontrolled entry into their ranks, as the model of the priestly dilemma in the previous section made clear; and this naturally implies that the migrant Brahmins would be careful to protect the welfare of their brethren "left behind" in their original homeland.

On these premises, the model introduced above for the Zoroastrian mission can readily be taken to describe the Brahmanical "mission", with $M$ now denoting not the members or adherents of the Zoroastrian religion but the active Brahmins and their patrons. There was a vast potential demand for the new religious and nonreligious services of this priestly class, while individual costs of joining would naturally be the higher, the farther away from the starting point they were in both geographical and cultural terms, consistent with the assumptions of the model. In this setting, it seems appropriate to posit that the Brahmin order's objective was to seek the most profitable expansion in their numbers and influence consistent with fully compensating the old members for their losses from the transition. After all, there is nothing in the technical machinery of the model that is specifically religious. All it takes is that the group functions as a producer cooperative engaging in benefit discrimination to maintain its unity and cohesiveness throughout its growth process, or else its market power would fall apart—and this seems like a perfect description of a self-perpetuating, self-congratulating *varna* like the Brahmins'.