1. (a) The following diagram demonstrates that the desired configuration can be achieved.

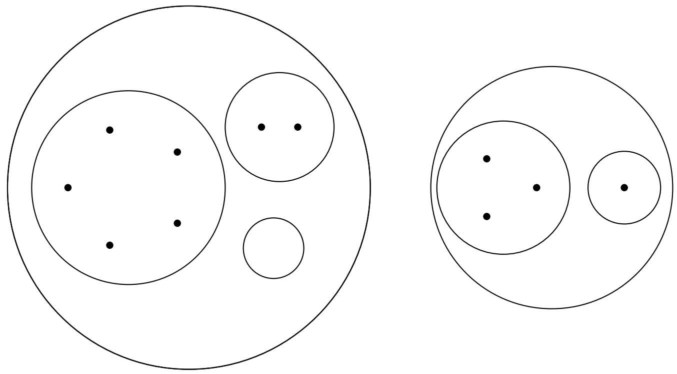

One way to come up with this diagram is by a counting argument. We will count the number of ordered pairs (C, D) where C is a circle and D is a dot contained inside of C. Suppose that each dot is enclosed in in k circles. Then the number of such pairs (C, D) is k for each circle, or $11k$. Alternatively, suppose that the number of dots contained in circle i is x_i, so then x_1, x_2, $\ldots$, x_7 must all be distinct. Then the number of pairs (C, D) is

$$11k = x_1 + x_2 + x_3 + x_4 + x_5 + x_6 + x_7.$$

Therefore, we want the sum of seven distinct numbers to be a multiple of 11. The smallest possible sum of seven distinct nonnegative integers is $0 + 1 + 2 + \cdots + 6 = 21$, so we try $k = 2$, with $\{x_1, x_2, \ldots, x_7\} = \{0, 1, 2, 3, 4, 5, 7\}$. In other words, we look for a diagram where each dot is contained in two circles, and the circles contain $0, 1, 2, 3, 4, 5$, and 7 dots. From here, we can come up with the diagram given in the beginning.

(b) For the sake of contradiction, assume that it is possible to draw five dots and five circles with these properties. By the same logic as in part (a), if each dot is enclosed in k circles, and if circle i contains x_i dots, then

$$5k = x_1 + x_2 + x_3 + x_4 + x_5.$$

We also know that the x_i are distinct nonnegative integers with $x_i \le 5$ (because a circle cannot contain more than 5 dots in it). Therefore,

$$10 = 0 + 1 + 2 + 3 + 4 \le x_1 + x_2 + x_3 + x_4 + x_5 \le 1 + 2 + 3 + 4 + 5 = 15,$$

which implies that $10 \le 5k \le 15$, so $k = 2$ or $k = 3$. We need to show both cases are impossible. If $k = 2$, then $x_1 + x_2 + x_3 + x_4 + x_5 = 10$, and the only way that we can choose the x_i to be distinct is if $\{x_1, x_2, x_3, x_4, x_5\} = \{0, 1, 2, 3, 4\}$. If $k = 3$, then $x_1 + x_2 + x_3 + x_4 + x_5 = 15$, and the only way that we can choose the x_i to be distinct with $0 \le x_i \le 5$ is if $\{x_1, x_2, x_3, x_4, x_5\} = \{1, 2, 3, 4, 5\}$. Either way, there exists circles with 4 dots, 3 dots, and 2 dots: call them C_4, C_3, and C_2. Note that C_3 must lie inside C_4, and C_2 must in turn lie inside C_3 (since there are only 2 dots outside of C_3, and they are separated by C_4). See the following diagram:

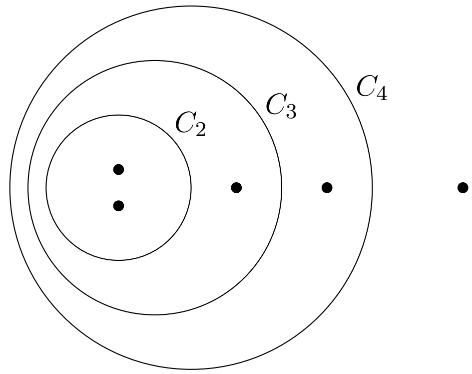

Now the point inside C_4 and outside C_3 has two fewer circles around it than the points inside C_2, so to make up, it must have at least two circles around it and outside of C_3. But then these two circles each contain only one dot, which is a contradiction. Therefore, it is impossible to draw a valid diagram with 5 circles and 5 dots. □

2. **Solution 1:** We can use Vieta's equations to find the coefficients of P and Q. Equating them, we get:

$$a+b+c = ab+bc+ca$$
$$ab+bc+ca = abc(a+b+c)$$
$$abc = (abc)^2$$

The last equation implies $abc = 0$ or 1. Since a, b, c are positive, $abc = 1$. By AM-GM,

$$\frac{a+b+c}{3} \geq \sqrt[3]{abc} = 1,$$

so $a+b+c \geq 3$. Conversely, 3 is achieved when $a = b = c = 1$, and in this case $P(x) = Q(x)$ so all constraints are satisfied. So the minimum possible value of $a+b+c$ given the constraints is $\boxed{3}$. □

Solution 2: Since polynomial factorization is unique, we know that the three roots a, b, c must be equal to ab, bc, ca (but possibly not in the same order). We first argue that one of the roots is 1. If $a = ab$, then since $a \neq 0$ we get $b = 1$; similarly we are done if $x = xy$ for any x, y. The only remaining case is if $a = bc$, $b = ca$, and $c = ab$. In this case, multiply the three equations together to get $abc = (abc)^2$, so $abc = 1$, and since $a = bc$, $a^2 = 1$, so $a = 1$.

Therefore, one of the roots is 1, and WLOG $a = 1$. Then $1, b, c$ must be equal to b, c, bc, which implies $bc = 1$. So the three roots must be $1, b, \frac{1}{b}$. We verify that as long as this holds, $P(x) = Q(x)$.

It remains to find the minimum possible value of $1 + b + \frac{1}{b}$. Since $b + \frac{1}{b}$ is minimized for $b = 1$, the minimum value is $1 + 1 + 1 = \boxed{3}$. □

3. We claim that Kim has a winning strategy. We demonstrate one possible strategy below. Kim's goal in this strategy is to guarantee the ability to place 4 tiles if necessary, one in each column, while at the same time limiting Li to placing at most 3 tiles. This allows Kim to make the final move.

   Kim begins by placing a tile in the middle of the second column. Then, no matter where Li places next, Kim places a tile in the third column, either in the top two squares or the bottom two squares – this is possible since Li's move can't have blocked both. Notice that at this point, two adjacent cells in the first column and two adjacent cells in the fourth column are free, such that Li can never place a rectangle in them; so Kim is guaranteed the ability to make two more moves, or four moves total. The diagram below shows the position of Kim's first two moves as blue tiles:

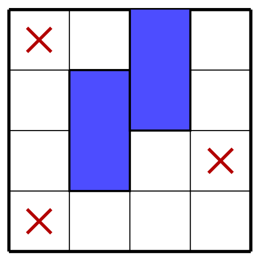

   For Kim's third move, consider three squares, marked as red Xs above: the top left corner, the bottom left corner, and the square in the middle two squares of the fourth column which is not adjacent to Kim's previously placed tile. Li has only placed two moves so far, so Li could not have covered all three Xs. For the third move, Kim places either in the first column or in the fourth column to block one of the Xs.

   After Kim's third move, there are still the four squares (in the first and fourth columns) that Li cannot place a rectangle on as noted above. But there is now one additional square (one of the Xs) that is blocked. Therefore, including Kim's initial moves, there are $4 + 4 + 1 = 9$ squares that Li cannot place a rectangle on in the course of this game, or rather $16 - 9 = 7$ squares that Li might be able to place a rectangle on. Thus Li can only place at most 3 tiles total. But Kim still has a fourth move available (either in the first or fourth column), so Kim wins. □

4. **Solution 1:** When reflected across the sides of the square, the four triangles' boundaries line up. In particular, the triangles' boundaries which are *inside* the square lie on only 3 line segments: a line ℓ from vertex B, a line m from vertex C, and a line n from vertex D. (The reflections of BX and BY lie on ℓ, the reflections of CY and CZ lie on m, and the reflections of DZ and DX lie on n. Since X must be below and to the right of A, the two reflections of AX lie outside of the square.)

   Orient the figure so that A is bottom-left, B top-left, C top-right, D bottom-right. Suppose towards a contradiction that there is a point not covered by the four reflected triangles. To not be covered by the reflection of XAB, it must be above ℓ. To not be covered by the reflection of YBC given that it is above ℓ, it must be below m. To not be covered by the reflection of ZCD given that it is below m, it must be below n. But if it is below n and inside the square then it is covered by the reflection of XDA. Contradiction, so no such point exists. □

Solution 2: For the sake of contradiction, assume that there exists a point P inside of square $ABCD$ such that P is not covered by any of the reflections of the four triangles. For the moment, we will focus only on square $ABCD$. Let P', P'', P''', and P'''' be the reflections of P across lines AB, BC, CD, and DA, respectively, as shown below.

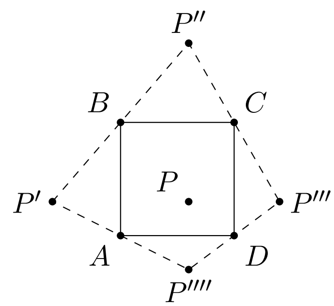

It can be seen that P', B, P'' are collinear, P'', C, P''' are collinear, P''', D, P'''' are collinear, and P'''', A, P' are collinear by comparing slopes. Now, since P is not covered by any of the four reflected triangles, P', P'', P''', and P'''' must all lie outside of triangle XYZ. Let r, s, and t be the lines extending the sides of triangle XYZ, such that r passes through B, s passes through C, and t passes through D. (In particular, X is the intersection of r and t, Y is the intersection of r and s, and Z is the intersection of s and t.)

Define "inside r" to mean below line r, "inside s" to mean below line s, and "inside t" to mean above line t. Then triangle XYZ consists of all points which are inside all of r, s, and t. Also, we may observe that P' is inside s and t, P'' is inside t, P''' is inside r, and P'''' is inside r and s. We summarize this with the following table:

	r	s	t
P'	?	inside	inside
P''	?	?	inside
P'''	inside	?	?
P''''	inside	inside	?

Because P' and P'' are collinear with B, which lies on r, one of them is outside r and one is inside r. Similarly, between P'' and P''', one is outside s and one is inside s. And between P''' and P'''', one is outside t and one is inside t. Subject to these constraints, it is impossible to fill out the above table without one row having *inside* in every column. This contradicts that P is not covered by any of the four reflections. □

5. **Solution 1:** We show that the number of ways is $12 \cdot 11 \cdot 10 = \boxed{1320}$. In fact, this generalizes. Suppose we want to color n points around a circle with a_1, a_2, a_3, ..., a_k of each color, where k is the number of colors and $a_i \geq 1$ for all i, satisfying the rule that each pair of colors is separated by a line. Then the number of ways is $n(n-1)(n-2)\cdots(n-k+1)$. While the argument presented here works in the general case, we focus on the specific case of 4 colors with $3, 3, 3$, and 3 of each color.

   To prove this, consider the following procedure for coloring the circle. We start out with three disks each of red, green and blue. For each of these three colors (excluding white), we pick a single point (called the source) where we place the stack of three disks of that color. Thus we have chosen three distinct points–a red source, a green source, and a blue source. Now, we color the circle as follows.

   - Start at point 1 and move around the points of the circle clockwise. As we move around the circle, we will be picking up and dropping off disks, so we will hold a vertical *stack* of disks in our hands (initially the stack is empty). As we visit each successive point, there are several possible scenarios.
     – If the point that we visit is a source that we have not visited yet, we pick up the three disks of that color and add those disks to the *top* of our stack. However, we also immediately place the top disk from our stack on that point.
     – If the point that we visit does not have any disks placed on it, and the stack of disks in our hand is empty, we move on to the next point.
     – If the point that we visit does not have any disks placed on it, and the stack of disks in our hand is not empty, we place the top disk from our stack on that point.
     – If the point that we visit already has a disk placed on it, move on to the next point.
   - We keep going around the points of the circle in clockwise order until we have visited each source, and then we continue going around the circle clockwise until our stack is gone.

   Since there are 9 disks and 12 points, this process will definitely stop. When it stops, we color the three points without a disk white.

   To illustrate this process, consider the following example. Suppose we initially place a red stack at point 2, a blue stack at point 4, and a green stack at point 6. As we move around the circle clockwise from point 1, our steps are demonstrated in the table below, where the "stack" refers to the stack of disks in our hand.

Point #	1	2	3	4	5	6	7	8	9	10	11	12
Stack before visiting point			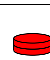	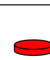	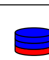	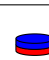			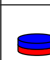	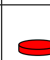		
Disks picked up												
Disks placed on current point		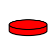	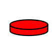	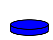	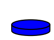	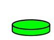	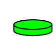	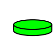	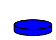	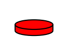		
Stack after visiting point		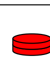	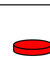	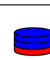	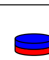		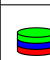	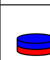	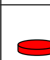			

Therefore, we can draw the following two diagrams, illustrating the initial placement of the stacks (left), and the final coloring that results (right).

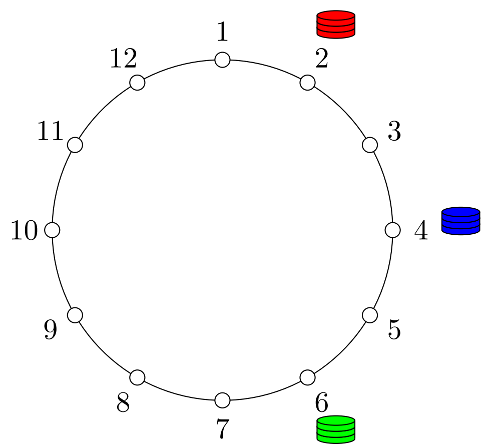

We first show this coloring is valid. As we use three disks of each color for red, green and blue disks, that leaves three points colored white, so we color three points of each color. Because of the stacking method, we never drop a red disk, then a blue disk, then a red disk, then a blue disk again; once we pick up the blue stack, we don't drop a red disk again again until the blue disks are gone. Also, we only color points with white on points where, every time crossing that point, we were holding no disks, so since we pick up a stack of a single color and drop all disks before we hold no disks again, we can draw a line that separates the disks of that color from the white points. Therefore, each pair of colors can be separated by a line.

Finally we have to show that this painting procedure generates every possible coloring exactly once. Starting from a coloring of the circle, the white dots separate the circle into three sections. Note that points of the same color can be in at most one of these sections, so for each color, we place a stack of three disks of that color on the first point in that section (in clockwise order) that is of that color. For example, in the above example, we split the circle into sections as shown below.

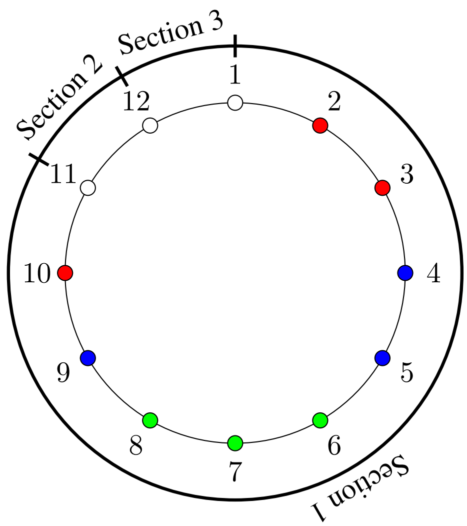

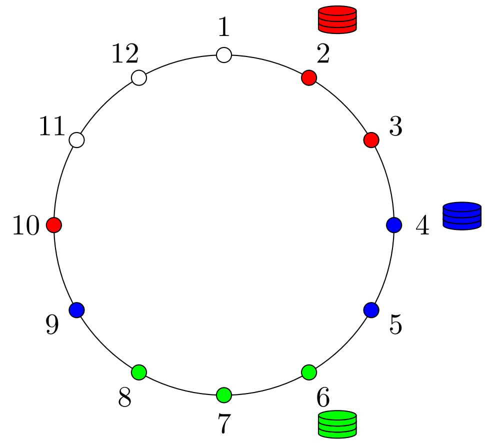

Since point 2 is the first red point in Section 1, we place the red stack at point 2. Since point 4 is the first blue point in Section 1, we place the blue stack at point 4. Since point 6 is the first green point in Section 1, we place the green stack at point 6.

If we do this and then do the previous procedure, we obtain the original coloring, and conversely, if we do the previous procedure and then this, we recover the original locations of the stacks. Therefore, the number of total colorings is the same as the number of ways to place the stacks. □

Solution 2: Suppose that we have a valid coloring of the 12 dots. Then we draw all segments between points of the same colors. This creates four triangles, which we call the red, blue, green, and white triangles, and we shade each triangle with their color as shown below.

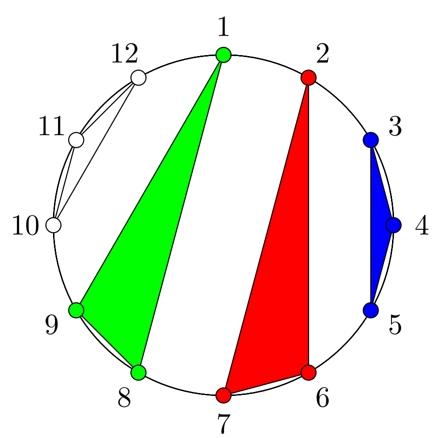

We first argue that a coloring is valid if and only if the four colored triangles do not overlap. If the coloring is valid, then consider for example the line separating the red and white points; then this line also separates the red and the white triangle. Conversely, if the red and white triangles do not overlap, then one of the three edges of the red triangle separates the red and white points (if it is moved a small amount away from the red triangle).

The 12 points split the circle into 12 equally-sized arcs. Call a triangle an (a, b, c) triangle if, when reading around the circle clockwise, the distance from the first point to the second is a, the distance from the second to the third is b, and the distance from the third back to the first is c (in number of arcs). For example, the blue triangle is a $(1, 1, 10)$-triangle. We claim that for each triangle, if it is an (a, b, c)-triangle, then $a + b + c = 12$ and $a, b, c \in \{1, 4, 7, 10\}$. The fact $a + b + c = 12$ is because the total distance around the circle is 12. For the second fact, note that since no two triangles intersect, each of the remaining three triangles must have all of its points on one of the three arcs that the triangle splits the circle into. Therefore, the number of points between two consecutive points on the initial triangle is always a multiple of 3, which means that the number of equally-sized arcs between those two consecutive points is one more than a multiple of 3. Hence $a, b, c \in \{1, 4, 7, 10\}$.

It follows that every triangle is either $(1, 4, 7)$, $(4, 4, 4)$, or $(1, 1, 10)$, up to reordering the triple. We proceed with casework based on which triples appear.

- **Case 1:** One of the triangles is $(4, 4, 4)$.
  There are 4 choices for the color of the triangle with triple $(4, 4, 4)$, and 4 rotations at which the $(4, 4, 4)$ triangle can be placed. Once the $(4, 4, 4)$ triangle is placed, the positions of the other three triangles are forced, so we simply have 3! ways to color the remaining triangles. Thus there are $4 \cdot 4 \cdot 3! = 96$ such colorings.
- **Case 2:** One of the triangles is $(1, 4, 7)$.
  The side of length 4 arcs forces a single $(1, 1, 10)$-triangle to be placed, whose position is forced. However, for the side of length 7 arcs, we have two subcases:
  - **Subcase 2.1:** The remaining two triangles are $(1, 1, 10)$ and $(1, 1, 10)$. In this case, there are 4 choices for the color of the $(1, 4, 7)$ triangle, and 24 choices for the position of the $(1, 4, 7)$ triangle (12 choices for the point between the arcs of length 4 and 7, and then 2 reflections). Since the remaining triangles are all $(1, 1, 10)$ triangles, their positions are uniquely determined, and there are 3! choices for their colors. Thus there are $24 \cdot 4 \cdot 3! = 576$ colorings here.
  - **Subcase 2.2:** The remaining two triangles are $(1, 4, 7)$ and $(1, 1, 10)$.
    In this case, note that the two $(1, 4, 7)$ triangles must have the sides of length 7 parallel and next to each other on either side of a diameter of the circle. There are 6 choices for this diameter. But we still must place the final point in each of these triangles, and due to reflections (perpendicular to the chosen axis), there are $2 \cdot 2$ choices for the third points of these triangles. Also, there are $4 \cdot 3$ choices for the colors of the two $(1, 4, 7)$ triangles. The positions of the final two triangles are uniquely determined, but there are 2 choices for the colors of the remaining triangles. Thus there are $6 \cdot (2 \cdot 2) \cdot (4 \cdot 3) \cdot 2 = 576$ colorings here.
- **Case 3:** All the triangles are $(1, 1, 10)$ triangles.
  In this case, the relative positions of the triangles are determined, but there are 3 distinct rotations. Then there are $4!$ ways to color the triangles, for a total of $3 \cdot 4! = 72$ such colorings.

Therefore, there are a total of $96 + 576 + 576 + 72 = \boxed{1320}$ colorings. □

6. We claim that the $\gcd$ is always 1 (so 1 is possible, and it is the only possibility). The proof will use the fact that p is prime, but will actually be valid for any positive integer $q \geq 2$.

   We use proof by contradiction: assume that the $\gcd$ is not 1; then there exists a prime number r, such that $r \mid p-1$ and $r \mid \frac{q^p-1}{q-1}$. We divide our proof into cases: either $r \mid q-1$, or $r \nmid q-1$.

   Case 1: $r \mid q-1$. In other words, $q \equiv 1 \pmod{r}$, so we have

   $$0 \equiv \frac{q^p-1}{q-1} = \sum_{i=0}^{p-1} q^i \equiv \sum_{i=0}^{p-1} 1^i = p = (p-1)+1 \equiv 1 \pmod{r},$$

   since $r \mid p-1$, which is a contradiction.

   Case 2: $r \nmid q-1$. Then $\frac{q^p-1}{q-1} \equiv 0 \pmod{r}$ if and only if $q^p - 1 \equiv 0 \pmod{r}$; that is, we can ignore the denominator. We can rewrite this as $q^p \equiv 1 \pmod{r}$. Let a be the smallest positive integer such that $q^a \equiv 1 \pmod{r}$ (this is often called the *order of* q *mod* r). Then $q^p \equiv 1 \pmod{r}$ if and only if $a \mid p$.

   On the other hand, Fermat's Little Theorem says that $q^{r-1} \equiv 1 \pmod{r}$, which implies $a \mid r-1$. And since $r \mid p-1$, $r-1 < r \leq p-1 < p$, so $a \mid r-1$ implies $a < p$.

   Therefore, $a \mid p$ and $a < p$. But p is prime, so $a = 1$. Finally, by definition of a, $q \equiv 1 \pmod{r}$, but that contradicts our assumption in Case 2 that $r \nmid q-1$. $\square$